VOLUME 12, NUMBER 8
JUNE 22, 1998

MICROPROCESSOR REPORT

THE INSIDERS' GUIDE TO MICROPROCESSOR HARDWARE

Merced Slips to Mid-2000

Delay Jeopardizes Attempt to Gain Performance Lead

by Linley Gwennap

Reality has reached out and tossed sand into the gears of Intel's product-development machinery. The company regretfully reported that an eight-month schedule slip has pushed the first volume shipments of its Merced processor from late 1999 to mid-2000. This slip will delay IA-64's penetration of the workstation and server markets and make it more difficult for Merced to achieve the performance lead, as Figure 1 shows. In the long term, however, the delay will probably have little effect on Intel's success in these markets.

Delay Caused by Poor Planning

Designing a new processor and a new instruction-set architecture (ISA) from scratch is always a long and daunting task. By pushing into the next decade, however, Intel's IA-64 effort is breaking two previous public commitments and is threatening to set a new record for gestation period.

The IA-64 effort was formally started in early 1994, when Intel and HP first began working together. The roots of this development effort stretch back even further, to research that HP and Intel had been separately conducting since 1991. When the partnership was announced (see MPR 6/20/94, p. 1), the partners said the first IA-64 products would ship before the end of the decade.

At the time, this statement seemed safe, as the internal plans were to complete the first IA-64 processor, code-named Merced, in 1998. Sources indicate that by 1996, however, the complexity of that chip was growing out of control, beyond what could be implemented in a 0.25-micron process. After examining and discarding proposals for a two-chip implementation, Intel decided to postpone Merced until its 0.18-micron process would be available. The new process allowed the design to be crammed onto a single chip. Unfortunately, this change delayed the ship date until mid-1999, when the new process would be ready.

Last fall, the company was confident enough in its progress to disclose the first details of the IA-64 instruction set (see MPR 10/27/97, p. 1) and reconfirm its commitment to 1999 shipments. At Microprocessor Forum, Intel's Fred Pollack promised that Merced would deliver "industry-leading performance" when it shipped. The design team at this point consisted of several hundred engineers, and the logical design was nearing completion.

Despite (or perhaps because of) this enormous staffing level, keeping the Merced program on schedule continued to be difficult. After a recent schedule review, senior management was shocked to discover that the chip was nowhere near tape out and in fact could not be expected to ship until the middle of 2000. Since Intel had publicly committed to 1999 shipments, it was forced to publicly acknowledge the change in plans.

Contrary to some rumors, the slip was not caused by delays in the 0.18-micron process; Intel still expects to ship 0.18-micron x86 products around 3Q99. Intel blames the latest delay in part on verification; it will take longer to test the design than originally thought. Given Intel's (and HP's)

Continued on page 6

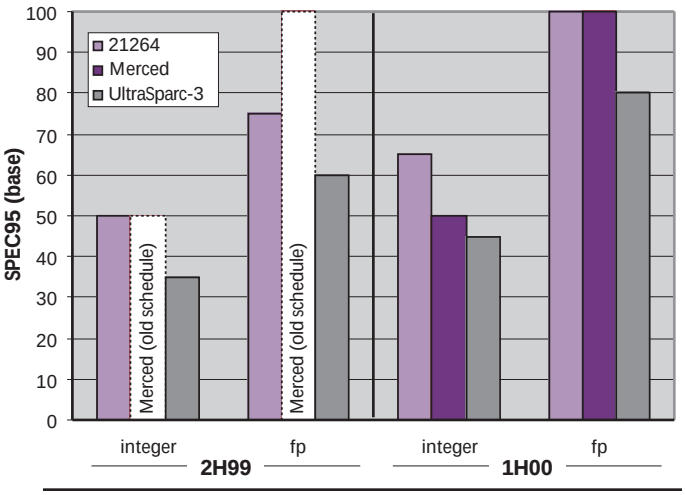

Figure 1. Merced had a good shot at gaining the performance lead in 2H99 on both integer and floating-point benchmarks, but by mid-2000 it could be merely among the pack in performance, particularly on the integer side. (Source: MDR estimates)

AT A GLANCE

MICROPROCESSOR REPORT

Founder and Executive Editor
Michael Slater
mslater@zd.com

Publisher and Editorial Director
Linley Gwennap
linley@zd.com

Editor in Chief
Keith Diefendorff
keithd@zd.com

Senior Editor
Jim Turley
jturley@zd.com

Senior PC Analyst
Peter N. Glaskowsky
png@zd.com

Editorial Assistant: Miriam Bauer

Editorial Board

Dennis Allison, Rich Belgard, Brian Case, Jeff Deutsch, Dave Epstein, Don Gaubatz, John Novitsky, Bernard Peuto, Nick Tredennick, John F. Wakerly

Editorial Office
298 S. Sunnyvale Avenue
Sunnyvale, CA 94086-6245
Phone: 408.328.3900 **Fax:** 408.737.2242

Microprocessor Report (ISSN 0899-9341) is published every three weeks, 17 issues per year. Rates are: N. America: $595 per year, $1,095 for two years. Europe: £450 per year, £795 for two years. Elsewhere: $695 per year, $1,295 for two years. Back issues are available.

Published by

MICRODESIGN RESOURCES

Business Office
874 Gravenstein Hwy. So., Suite 14
Sebastopol, CA 95472
Phone: 707.824.4004 **Fax:** 707.823.0504
Subscriptions: 707.824.4001
cs@mdr.zd.com

World Wide Web: *www.MDRonline.com*

Computer Press Award, Best Newsletter, Winner, 1993, 1994; Runner-Up, 1996

Printed on recycled paper with soy ink.

THE EDITORIAL VIEW

Playing the Digital Convergence Game

WebTV Could Be Converged Out of Existence

When trying to understand politics, the rule is to follow the money. In considering the opportunities for digital convergence, I have a new rule: follow the MIPS. In general, a function that requires fewer processor MIPS and a simpler feature set can be subsumed into a device that provides more MIPS and a more complex feature set. This rule spells trouble for the current WebTV, among others.

The devices that connect to a television are ripe for digital convergence. These devices are generically called set-top boxes, although it's easy to have so many that they end up on top of, underneath, and somewhere near the TV. Today's leading-edge consumer might have a DVD player, DSS satellite receiver or cable decoder box, video-game system, and WebTV-type device. Although few people today have all of these products, they are all increasing in popularity.

The total cost of all these units is large, and having several discrete boxes makes the combination difficult to use. Simply cabling all these units together and finding a place to put them can be difficult. Grouping some functions would reduce overall cost and simplify the picture.

Let's follow the MIPS. The DVD player combines a spindle, a microcontroller, and an MPEG decoder. The DSS unit is chock-full of electronics, with an analog front end, digital signal processing, and an MPEG decoder as well as a phone-line interface. The video game has a high-performance CPU, a 3D graphics chip, and, if it is a Sega Saturn or Sony PlayStation, a CD-ROM spindle. A WebTV unit has a similar CPU and uses a phone line to connect to the Internet. All of these devices have a video output, of course.

The WebTV looks like the first to disappear. Sega has already announced that its next-generation device, called Dreamcast (see MPR 6/1/98, p. 8), will include a built-in modem for multiplayer gaming. The same modem can be used to surf the Web, and Dreamcast will include browser software for this purpose. Like the current Sega Saturn, Dreamcast uses a CD-ROM to access its game software.

Additional convergence is possible. The SH-4 processor in the Dreamcast system generates enough MIPS for MPEG-2 decoding. As the cost of DVD drives continues to drop and approaches that of CD-ROM drives, a slightly modified Dreamcast with the right software could play DVD movies as well as video games, replacing the DVD player.

The makers of DVD players, of course, don't want to be left out of the action. Consumer-electronics giants Thomson and Toshiba are working with a small semiconductor maker, VM Labs (see MPR 6/22/98, p. 22), to add video-game functions to their DVD players. The VM Labs media processor has enough performance for MPEG-2 decoding as well as the 3D graphics needed for popular video games.

This chip could also be used to replace the hardwired MPEG-2 decoder in a DSS receiver or digital cable box. In this way, these systems could also be converted to video-game players. One drawback to this combination is that a DSS or cable box cannot read CD-ROMs, which are used to distribute many video games today. A CD-ROM could be added, of course, or the system could accept games stored on ROM cartridges (as in the Nintendo 64) or even download software from the Internet.

Going whole hog and adding a DVD drive to the DSS box would allow an all-in-one system that could access satellite programming, play DVD movies, play video games, and surf the Web. Such a device would be the lowest-cost solution for this combination of features, and it should be easier to install and use than a set of discrete units.

The problem with integration is that not everyone wants all of these functions. Some don't care about video games. Others don't need satellite television. In addition, consumers have different needs and desires even for the same function. High-end DVD players offer a range of features not found in basic DVD players. Preferences are even stronger in the video-game space, where the software and capabilities of each platform vary widely.

Therefore, there will continue to be a variety of standalone products in this space, particularly products that provide some high-end features. Many consumers will want to pick a specific DVD player and combine it with an existing video game, for instance. But consumers looking for basic functions may be served well by an integrated unit.

Web browsing is a basic function that can be easily integrated into video games, DSS receivers, and even televisions themselves. I expect such integration to become widespread, reducing the need for standalone WebTV-type devices.

Where games are desired, the video-game companies are in the driver's seat. A DVD manufacturer can't add Nintendo compatibility to its systems, but Nintendo can easily add a basic DVD playback function or even a DSS interface. Consumers who want video games should be able to get these other capabilities at little or no extra cost. If the video-game makers don't seize this opportunity, however, others will. M

IBM Abandons Somerset

The initial era of PowerPC has come to a close as Motorola and IBM agreed to part ways on the Somerset Design Center. Motorola takes ownership of the design center and all future designs, including the G4. IBM will continue to develop its Power processor family, which is mainly aimed at servers, but has no current plans to build or sell the G4. Within two years, we expect this change will leave Apple solely dependent on Motorola's chips.

Somerset was the centerpiece of the original PowerPC agreement (see MPR 10/16/91, p. 1) between the two companies and Apple; jointly staffed by IBM and Motorola, Somerset was tasked to simultaneously design a family of PowerPC processors, including the 601, 603, 604, and 620. After turning out the first three chips with impressive alacrity, Somerset whiffed on the 620. Worse, development of a second generation of cores lagged; seven years after its founding, Somerset has yet to produce a true second-generation part.

As in Arthurian legend, Somerset was the start of a Holy quest: to wrest a significant chunk of the desktop PC market away from Intel. At the time, Apple held 10% of that market, and IBM held about the same. Since then, however, Apple's market share has dwindled to a mere 4%, and IBM's PC group refused to accept the PowerPC chips.

This situation leaves too little volume to support two sources of PowerPC chips. With IBM maintaining its own line of Power processors, it was getting little from Somerset other than a share of Apple's shrinking business. As we pointed out last fall (see MPR 10/6/97, p. 3), Somerset had become superfluous.

IBM will continue to market RS/6000 workstations and servers based on the 604 as well as its Power chips. The 604 suffers in the workstation market due to its 32-bit ISA and modest FP performance. IBM hopes to solve these problems by bringing the Power3 (see MPR 11/17/97, p. 23) into the volume workstation market in 1999. IBM has also introduced a line of Pentium II workstations and is likely to market Merced workstations when that processor is available in 2000 (see MPR 6/22/98, p. 1). These Merced systems will complement IBM's RS/6000 workstation line.

IBM will continue selling processors to Apple in the near term, but without access to the G4, this business will eventually fade away. Ironically, one of Apple's key goals in creating the PowerPC partnership was to create a dual-sourced architecture. By 2000, Apple will be back where it was in the 1980s: solely dependent on Motorola for its CPUs.

Both Motorola and IBM will continue to develop their own PowerPC chips for embedded applications (see MPR 6/22/98, p. 10). But despite their parting on Somerset, they will continue to work together on Book E (see MPR 10/27/97, p. 10), which will define extensions to the current architecture for future embedded processors.

In summary, the future of PowerPC has fractured. IBM will continue to develop its own server chips, which it does not sell openly. Motorola, through its Somerset design team, will develop processors for Apple. And both companies will separately develop embedded PowerPC chips, ensuring compatibility via Book E. The new structure provides a flow of processors for the few remaining system makers, along with the opportunity for both PowerPC chip makers to exploit emerging embedded opportunities. —*L.G.*

New Intel Competitor Rises

Rise Technology *(www.rise.com)*, a new fabless semiconductor startup, has disclosed plans for its first microprocessor. The mP6 is an x86-compatible processor targeted at the sub-$1,000 PC market. With this focus, it will compete head-to-head with IDT's WinChip and Cyrix's highly integrated processors.

The company was founded in 1993 by David Lin and has since grown to more than 80 people. A second division is located in Taiwan, close to many high-volume PC manufacturers. Rise is privately funded by undisclosed first-tier venture capitalists, investment bankers, and PC companies.

Rise's strategy is to zero in on small niches—of a few million units—in the PC industry, where Intel is not focused and where Rise can establish a competitive advantage. Its first target is the low-cost noncorporate notebook PC market. This segment is quite small, as only 20% of all PCs are notebooks, and most of those go to corporate customers, who strongly prefer Intel processors. But Rise hopes this loyalty will erode, as it appears to be doing in the consumer segment.

No technical details about the mP6 will be disclosed until October's Microprocessor Forum, but Rise claims the chip will have Pentium II-class performance and the highest per-clock MMX performance of any processor available. Sources indicate that mP6 does not implement the 3DNow extensions (see MPR 6/1/98, p. 18). Advanced power-reduction techniques were employed to bring the mP6 within the power envelope of a notebook PC.

The mP6 will also be suitable for use in low-cost desktop PCs. Rise will offer it in Socket 7, allowing the chip to compete directly with other low-cost desktop processors. The mP6 is bucking the integration trend pursued by Cyrix and IDT. Rise says it may develop integrated parts later, and it is planning to license its technology to selected partners that could produce such products.

Rise believes its processor does not infringe any Intel patents, but it is using an Intel-licensed fab to be safe. The startup did not disclose its foundry partner, but possible candidates include IBM, Texas Instruments, and SGS-Thomson.

The mP6 will be available in PGA and BGA packages by the end of the year and is expected to sample this summer. Chip-set vendors Acer, VIA, and Utron have announced they

are working with Rise to ensure compatibility between their Socket 7 chip sets and the Rise processor. No pricing for the mP6 has been announced.

If Rise delivers Pentium II performance within the notebook power envelope (about 8 W), it will be competitive in its chosen niche. On the desktop, however, it will have to compete mainly on price. Rise did not disclose the mP6's die size, but the company will be hard-pressed to match IDT's 58-mm^2 WinChip 2 3D (see MPR 6/1/98, p. 1), its chief competitor in this market. AMD and Cyrix processors are also expected to join the fight for the low-cost desktop market.

IDT has already established its x86 compatibility, which Rise must yet do, and aggressive pursuit of its WinChip roadmap could make life tough for the new vendor. Fortunately for Rise, IDT cannot supply the entire market, which should leave the startup a few openings. —*K.D.*

Intel Accelerates Price Cuts, Katmai, Celeron

Faced with sagging PC demand and increased competition, Intel has responded by accelerating price cuts on its desktop processors by seven weeks. Additional price cuts are likely next month as well, resulting in a large total decrease. The company also pulled in the schedules of its high-end Katmai and low-end Celeron CPUs.

Intel's first Celeron product, the 266-MHz Covington (see MPR 4/20/98, p. 14) did not meet with rave reviews, to put it politely. At 266 MHz, the cacheless Covington delivers lower performance than a 233-MHz Pentium/MMX on PC application benchmarks such as Winstone 98. To improve Celeron's acceptance, Intel has officially added to the mix a 300-MHz version of Covington with immediate availability. Since Covington uses the same Deschutes CPU as Pentium II, producing the higher clock speed is no problem.

The new part, sold as the Celeron-300, delivers better performance than the fastest desktop Pentium/MMX on most benchmarks, making it a more viable option for low-end PCs. Both Celeron products are now shipping, and vendors such as Compaq, Dell, HP, and IBM have announced Celeron-based PCs.

To further compensate for Covington's poor performance, Intel has cut the price of the Celeron-266 by 32% almost before the part began shipping. The 266-MHz processor now lists for $106, the same price as the 233-MHz Pentium/MMX. The Celeron-300 was introduced at $159, filling the price point of the initial Celeron.

Intel's forthcoming Mendocino chip will add a 128K on-chip L2 cache, greatly improving performance. The company plans to deploy 300- and 333-MHz versions of this chip in 4Q98. To distinguish the 300-MHz Mendocino from the same-speed Covington, the former will be marketed as the Celeron-300a. The faster Mendocinos will not use the "a" designation, and the whole issue should quickly become moot as the weaker Covington fades from the market.

While revising its Celeron prices, Intel also trimmed the 1,000-piece list prices of its other desktop chips in an unusual midquarter price cut. As the table below shows, the price cuts, effective June 7, ranged from 16% to 20%, except for the high-end 400-MHz product, which lost only 12%. Intel cut the price of the Pentium II-233 to just $161; PC makers that are unhappy with Celeron can choose this part for only $2 more.

Although these cuts are smaller than Intel's typical quarterly moves, the company plans to cut prices further at the end of July, its normal schedule for price changes. As a result, the total desktop price cuts for this quarter are likely to be significantly larger than usual. The latest price changes do not affect Intel's mobile processors.

Processor	4/15/98	6/7/98	%CHG
Pentium II-400	$824	$722	–12%
Pentium II-350	$621	$519	–16%
Pentium II-333	$492	$412	–16%
Pentium II-300	$375	$305	–19%
Pentium II-266	$246	$198	–20%
Pentium II-233	$198	$161	–19%
Celeron-300	n/a	$159	n/a
Celeron-266	$155	$106	–32%
Pentium/MMX-233	$134	$106	–21%
Pentium/MMX-200	$95	$95	0%

Intel further disclosed that the first silicon of its Katmai processor is looking good enough that the company now expects to ship the part in 1Q99, a bit earlier than previously scheduled. This change will help reduce the gap between Intel's Pentium II-450, expected to ship in 3Q98, and Katmai, avoiding a long plateau in high-end performance.

The company said Katmai will ship initially at clock speeds of up to 500 MHz. This extra speed, along with some expected minor core changes, will allow it to offer a small performance advantage over Pentium II on typical PC applications and a large advantage on applications that use the multimedia extensions known as Katmai new instructions.

Overall, these product changes address weaknesses in Intel's low-end strategy. The Covington-300 should improve Celeron's acceptance, and the plan to replace that part with the Mendocino-based "300a" should quickly sweep Covington under the rug by early 1999. With Katmai and Mendocino coming on strong at that time, Intel should be well positioned in 1H99. —*L.G.*

Ross Technology Prepares for Shutdown

Facing deepening financial problems, Ross Technology *(www.ross.com)* has laid off 40% of its workforce and placed all products on "end of life" status. The company lost more than $40 million over the past two years and reports that sales of its 32-bit SPARC processors are "deteriorating." During the past several years, most SPARC system makers have abandoned the architecture in favor of the x86 family, limiting sales opportunities.

Founded by Roger Ross (who led the development of Motorola's 88000) in 1988, the eponymous company began

Continued on page 21

Merced

Continued from page 1

long experience in designing new processors and even new ISAs, however, this lack of foresight seems improbable. Other problems may exist with the design, but Intel is not willing to discuss them.

Competitive Performance May Suffer

Intel denies rumors that Merced is having trouble meeting its clock-speed and performance goals. Even if its performance is still on target, however, Merced will have more of a problem exceeding the competition than it would have had in 1999. Based on Moore's Law (a concept Intel should be familiar with), competitive performance should increase by about 25% during a six-month slip.

Our estimates support this assertion. We expect the fastest processor in 1999 will be the 0.25-micron 21264, delivering an estimated 50 SPECint95 and 75 SPECfp95, as Figure 1 shows. Had Merced been able to ship in that year, we believe it would have exceeded the Alpha chip's floating-point performance and potentially its integer performance as well, fulfilling Pollack's promise.

Part of Merced's performance comes from its use of 0.18-micron process technology, giving it an advantage over 0.25-micron processors. Under the previous plan, Merced was to be one of the first products off Intel's—or anyone else's—0.18-micron process. By mid-2000, however, 0.18-micron technology will be in common use for leading-edge microprocessors.

In particular, Intel is committed to delivering a 0.18-micron version of the 21264 processor to Digital as part of the two companies' foundry agreement (see MPR 11/17/97, p. 1). We expect that chip to reach 65 SPECint95 and 100 SPECfp95. Similarly, Sun's UltraSparc-3 (see MPR 10/27/97, p. 29) should also ship in a 0.18-micron process by that time, achieving 45 SPECint95 and 80 SPECfp95.

Since the schedule slip does little to improve Merced's performance, we expect it will have a tough time exceeding the integer performance of the 0.18-micron 21264. And its performance advantage over processors such as UltraSparc-3 and IBM's Power3 will be narrowed. Thus, Intel may not be able to claim bragging rights for its new architecture.

Sources indicate that Intel and HP are now pinning their performance hopes on the second-generation IA-64 processor, code-named McKinley. Intel had previously committed to delivering this device in 2001 with twice the performance of Merced in the same CMOS process. The company says the Merced slip will not affect McKinley, which is being developed by a separate design team.

Market Impact Delayed, Not Diminished

Even if Merced is not the fastest microprocessor in the world, its impact on the market won't be significantly diminished. The list of system vendors that have already committed to IA-64 is long and growing longer. It already includes Compaq, Data General, Dell, Digital, Groupe Bull, HP, Hitachi, IBM, ICL, Micron, NCR, NEC, Sequent, Siemens-Nixdorf, Silicon Graphics, Stratus, and Unisys. None of these companies is likely to reconsider its support of IA-64 because of this slip. With the backing of these companies, IA-64 has a clear path to taking over the majority of today's RISC-based workstation and server markets.

Of course, IA-64's impact will be delayed, and these vendors must make some tactical adjustments. Most base their current workstations and servers on Intel's Pentium Pro and Pentium II processors. These vendors will have to stay with the x86 line a bit longer before moving to Merced. Tanner (see MPR 3/9/98, p. 4), an x86 processor that plugs into the same Slot M as Merced, will allow system makers to develop and deploy Slot M systems that can later be upgraded to Merced. Given the latest slip, we expect Intel will develop a version of its Willamette x86 processor that also plugs into Slot M. This product could extend the company's x86 server line in 1H00.

Other vendors, such as HP and SGI, are moving from RISC to IA-64. HP is developing a processor called the PA-8700 to extend its current RISC line, tiding over customers until Merced systems are available. HP claims at least some of its PA-8700 systems will be upgradable to Merced. Similarly, SGI will rely on its forthcoming R14000 until Merced ships and will also offer upgradable systems. The delay will extend the current period of uncertainty for these vendors' customers, making them more vulnerable to poaching from non-IA-64 vendors such as Sun. But both HP and SGI are prepared to ride out the IA-64 delay.

A conspiracy theorist might claim Intel had this plan all along: prematurely announce the part to drive all competition from the market, then disclose the real schedule. Our sources indicate this schedule slip was a surprise at all levels of the organization, but the ultimate outcome is the same as in the conspiracy theory.

Financial Impact Minimal

The announcement of Merced's slip caused Intel's stock price to drop by 8%, but the financial markets seem to have overestimated Merced's impact on Intel's revenue. We had estimated Merced's 1999 revenue to be $400 million, less than 2% of the company's annual $25 billion revenue. With the new schedule, IA-64 is likely to contribute only a few percent to Intel's top line in 2000 and 5–10% in 2001. Not bad business if you can get it, but hardly enough to affect current stock prices.

Thus, the Merced delay is embarrassing for Intel but survivable. The company has much more pressing problems, like avoiding antitrust sanctions (see MPR 6/22/98, p. 8) and finding compelling applications that require the performance of a 400-MHz Pentium II, much less an 800-MHz Merced. Intel needs to get its scheduling problems under control, however, to prevent further

CMOS Image Sensors Challenge CCDs

CMOS Makers in Search of the Holy Grail—the Digital Camera on a Chip

by Keith Diefendorff

Today, charge-coupled devices (CCDs) fill most of the sockets for digital image sensors in all types of devices. But CCDs have characteristics that make them less than ideal for inexpensive handheld devices. Several semiconductor manufacturers hope to exploit the power, size, and cost advantages of CMOS image sensors to capture the exploding market for sensors in digital video and still cameras.

Unlike CCDs, CMOS image sensors are built in standard CMOS semiconductor processes that are amenable to integrating analog and digital signal-processing circuits on the same die as the sensor array. In the past, CMOS image sensors have not been able to deliver image quality comparable to that of CCDs. But new developments, especially active-pixel-sensor technology, eliminate this shortcoming.

Issues remain, but it is now apparent that CMOS image sensors will displace CCDs in consumer-grade digital cameras. Fax machines, scanners, automobile-vision systems, and other devices will likely go the same way, as the issues are similar. In this article, however, we focus primarily on the most interesting opportunity for semiconductor manufacturers—the digital-camera market. Future Image and IDC, in their *Digital Camera Market Review and Forecast*, estimate the worldwide market for digital still cameras alone will be 8.5 million units by 2001.

Recognizing a new business opportunity in the digital-imaging market and the potential for it to boost demand for high-performance PC microprocessors, Intel has established a digital-video and imaging division that will market digital cameras for PCs using CMOS image sensors that Intel itself manufactures. TI, Motorola, Toshiba, Rockwell, and others also have their sights set on this emerging market.

Converting Photons Into Bits

Figure 1 shows the process of going from an optical image to a digital one. First, the optical image is focused on a digital image sensor. The image sensor comprises a rectangular matrix of light-sensitive elements, each of which represent one pixel of the image.

The number of pixels in the sensor array determines the resolution of the final image. On the low end, arrays of 320 × 240 pixels are adequate for NTSC video. Arrays of 640 × 480 are used for low-cost digital cameras and Web publishing. High-end cameras need at least 1,000 × 800 resolution, but these are already moving into the mainstream. The attention is now on 1,300 × 1,000 arrays, which begin to make digital cameras interesting alternatives to film-based snapshot cameras in the consumer market.

Each light sensor in the array is a photodetector that converts photons impinging on it to an electric charge via the photoelectric effect. The charge is integrated over a period of time long enough to collect a detectable amount of charge but short enough to avoid saturating the storage elements. This period is analogous to film exposure time. In general, photodetectors are more sensitive to light than film, which allows digital cameras to have short exposure times ranging from $^1/_{50}$ to $^1/_{10,000}$ of a second.

After the charge has been collected, it is transferred out of the array, where it is converted to a voltage proportional to the magnitude of the charge. Analog signal-conditioning circuits amplify and filter the signal, which is weak (a few millivolts in low light) and susceptible to noise injection. The signal-to-noise ratio (SNR) at this point is extremely important to the quality of the final image. A 42-dB SNRs is needed for a VHS-quality still picture. Typical CCD sensors today achieve SNRs better than 60 dB.

After conditioning, the analog signals are digitized by an analog-to-digital converter (ADC) into binary values that are proportional to the intensity of the light impinging on the respective sensors. The ADCs usually have a range of 8 to 10 bits, to cover the entire dynamic range of the signal. Subsequent processing is performed in the digital domain.

Unfortunately, photodetectors are monochromatic, so the digital image at this point is only grayscale. For creating a color image, a popular method is to print a polyimide mask of color filters over sensors in the array. A mosaic of red,

Continued on page 12

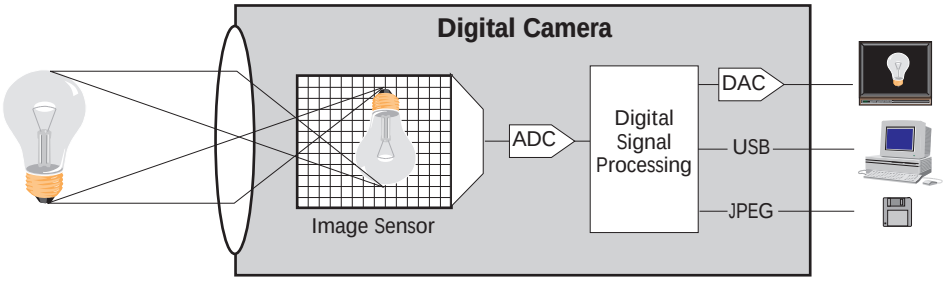

Figure 1. In a digital camera, the image is focused on an image sensor made up of an array of light-sensitive elements (pixels). Analog signals corresponding to the intensity of light falling on each pixel are converted to digital values by the ADC and processed into a final color image by digital-signal-processing circuits.

FTC Files Suit Against Intel

U.S. Government Challenges Intel Business Practices

by Linley Gwennap

Attempting to define the appropriate business practices of a company with monopoly power, the U.S. Federal Trade Commission (FTC) has filed a suit against microprocessor giant Intel, alleging anticompetitive behavior that is illegal under federal law. Intel quickly responded that it believes its actions have been consistent with existing case law, and that the FTC is creating "a new legal theory" to deal with this case.

If the FTC prevails in its case, it would clarify the responsibilities of technology companies—such as Intel, Microsoft, and Cisco—with dominant market shares. The government asserts that such companies cannot cut off the supply of information and products to customers unless the action serves a "legitimate, procompetitive purpose." In at least three situations, Intel took actions that the FTC deems anticompetitive.

While Intel will fight the case, the penalty for losing is not severe. The FTC is asking only that Intel "cease and desist from directly or indirectly discriminating" against its customers with regard to supplying advance product information, prototypes, and actual products. We doubt that such a course of action would cause Intel to lose a significant amount of business to its competitors.

Intel Freely Admits Actions

This case is unusual in that Intel freely admits to most of the allegedly illegal actions. The FTC has built its filing around three interactions between Intel and its customers. In the first case, Digital lost access to Intel's technical information and prototype processors after it sued Intel for patent infringement (see MPR 6/2/97, p. 16). Had Digital not come to a settlement with Intel (see MPR 11/17/97, p. 1), this situation could have prevented Digital from developing viable systems that used Intel chips.

Intel took similar action against Intergraph when that company attempted to assert its own patents against some of Intel's customers (see MPR 12/8/97, p. 4). Like Digital, Intergraph couldn't develop new products until Intel was forced by a federal judge (see MPR 5/11/98, p. 16) to supply the necessary information.

The FTC's third example involves Compaq, which sued Packard Bell, another PC maker, in 1994 in a dispute over motherboard patents. Packard Bell, which was using motherboards purchased from Intel, asked Intel to indemnify, or protect, it from Compaq's suit. Intel then went to Compaq seeking rights to Compaq's motherboard patents and, when the PC maker refused to cede them, cut off its access to technical information.

Compaq extended its lawsuit to cover Intel as well as Packard Bell, but the companies eventually came to a settlement. Like Digital, Compaq acceded to Intel's demands, licensing its patents rather than attempting to compete without Intel's support.

Intel doesn't deny the basic facts that it withdrew from these companies access to technical information on unreleased products (books with yellow, orange, or red covers, collectively known as "color books") and prototypes of these products. The color books and prototypes are controlled under nondisclosure agreements (NDAs), so Intel simply terminated these agreements. The company says it has "the legal right to assert its intellectual property rights as a defense to an attack on its core microprocessor business."

In the Digital case, Intel responded to a surprise lawsuit that requested Intel to stop shipping most of its products. In the Intergraph and Compaq cases, however, Intel terminated the NDAs before those companies filed suit. Compaq never threatened Intel's core microprocessor business; its suit was focused on motherboard issues.

FTC Stretches Antitrust Doctrine

In any case, the FTC believes Intel had no right to terminate NDAs in these circumstances. Although Intel will not admit to being a monopolist, the company's market share is universally agreed to be 80% or more, well within the range that is generally considered to be a monopoly. In issuing an injunction in the Intergraph case, Judge Edwin Nelson found that "Intel has monopoly power in the relevant market of high-performance CPUs."

Market share is not the only criterion in being declared a monopolist. The ability to set prices and to control the market are key factors. Intel's lawyers claim the company's recent price cuts (see MPR 6/22/98, p. 5) are a response to competition, although in other public forums the company rarely mentions competition as a reason for cutting its prices.

Intel's control over the market is frequently demonstrated, most recently by PC makers' transition from Socket 7 to Slot 1, which is occurring rapidly despite opposition from Intel's competitors and even some of its customers. This transition also leaves Intel open to charges that it has a monopoly in Slot 1 processors, due to its 100% share in this market segment.

It is not illegal to be a monopolist. But under the Sherman Antitrust Act, the Federal Trade Commission Act, and subsequent case law, it is illegal to use monopoly power to entrench a monopoly or to restrain competition. By using its monopoly power to extort patent licenses from its customers, Intel is acting illegally, says the FTC.

Case law, however, is not clear on this point. Antitrust law has generally been used to prevent monopolists from harming their competitors, not their customers. Proving restraint of trade in this case may be difficult.

One argument is that Intel's actions reduce competition in the systems market by preventing companies such as Intergraph from building viable products. Intel points out, however, that eliminating one of many system vendors will not perceptibly reduce the overall level of system competition. Although this treatment is harsh on a particular system maker, the impact on system buyers is negligible.

The FTC filing makes passing reference to the harm done to the targeted system vendors, but its key argument is different. The filing asserts that Intel's coercive actions will "diminish the incentives of the industry to develop new and improved microprocessor and related technologies." Thus, this conduct "reduces competition to develop … future generations of microprocessor products."

This argument seems odd, however, since neither Compaq nor Intergraph were developing microprocessor products that might compete with Intel's. Digital's Alpha processors do compete with Intel's chips, but Intel's actions threatened Digital's x86-based business, not its Alpha-based business. Compaq and Intel compete in the motherboard arena, but the FTC focuses on the microprocessor market.

A third argument is that Intel gains a market advantage over its real competitors, such as AMD and Cyrix, by obtaining patent rights from its customers in a manner that its smaller rivals can't duplicate. In this way, Intel can gain access to more intellectual property than its competitors can. The FTC filing, however, does not include this argument.

First Step Is Administrative Proceeding

After voting 3-1 to pursue the suit, the FTC now turns the case over to an FTC administrative law judge, making the case an internal matter. Although the administrative judge performs many of the same functions as a federal judge, the proceeding is simpler than a full trial and is not allowed to last more than 12 months. At the completion of the proceeding, the judge will make a recommendation to the FTC commissioners, who will then vote on whether to accept it.

If the FTC votes to sanction Intel, it can order "necessary and appropriate relief." Its filing outlines specific requests, mainly that Intel must provide color books, prototypes, and products in a similar fashion to "similarly situated" customers. It requests that Intel not cut off any customer without "a legitimate business consideration." Issues of intellectual-property disputes or competitive positions of the customer are not considered legitimate reasons to stop serving a customer.

Intel says it will contest the FTC's action through the administrative proceeding and ultimately appeal its decision to a federal court, if necessary. Although the initial proceeding must be completed by the middle of 1999, an appeal could stretch out the timing of any sanctions.

For More Information

For a copy of the FTC's filing against Intel, access the Web at *www.ftc.gov/os/9806/intelfin.cmp.htm*.

New Legal Interpretations Required

In deciding this case, the administrative judge (and potentially a federal appeals judge) has a difficult task. As noted, most of the basic facts are not in dispute, only the legal theory. In Intel's favor, existing antitrust precedents do not seem to apply here. Antitrust experts, even those who favor the FTC's position, could not point to a previous case where a company was sanctioned for mistreating customers. Although Intel's actions have clearly not been "procompetitive," that does not necessarily make them anticompetitive.

The FTC did not raise the "essential facility" argument in its filing, but this concept could arise during the proceeding. In the Intergraph hearing, Judge Nelson found that Intel's processors are essential to the viability of its customers, just as electricity and phone service are essential.

The classic case of an essential facility is a railroad bridge that was the only way across the Mississippi River. The railroad company that built it was forced to provide access to its competitors, because the bridge was an essential facility. In this case, the law forced an increase in direct competition. It did not affect the railroad's dealings with its customers, which would make it more applicable to the Intel case.

The basic antitrust laws, however, are very broad. The Sherman Act is only two paragraphs long, and its essence is to prevent a monopolist from restraining competition. The Act does not describe what constitutes a restraint of competition; this vagueness gives judges much leeway in interpreting the law.

The judge, like many observers, may take a dim view of Intel's heavy-handed tactics. Although removing one of many competitors from the PC or workstation market may do little to reduce overall competition, giving Intel the power to destroy another company's business seems unfair. A judge who feels Intel has been unfair, however, must still find a defensible legal precedent on which to base a ruling.

Intel May Choose to Settle

Although it might prevail in court, Intel is likely to settle this case, potentially after the initial proceeding but before the final decision. The cost to Intel of modifying its business practices is small, and a settlement could allow it to make "voluntary" changes to its business practices.

If Intel is found to be a monopolist, however, the finding would open the door to further lawsuits from the FTC and from Intel's competitors. These suits could attack Intel's moves in secondary markets such as system logic, motherboards, and graphics. The potential for onerous restrictions in these areas is a danger Intel can't afford to face. M

Motorola and Lucent to Share DSPs

Motorola and Lucent Technologies have decided to exchange DSP architectures and collaborate on the development of new ones. The first new cores to spring from this alliance are due in 1999 and will position both companies to compete more strongly against DSP dominator Texas Instruments.

The next-generation DSP cores will be designed at a jointly owned design center in Atlanta called Star Core, which will open in 3Q98. The DSP architecture will be fully defined, and the first core implementation will be ready for use, by mid-1999. Chips based on the new architecture are expected from both companies before the end of that year.

Defining a new DSP and designing a core implementation generally takes much longer than 12 months, and this is exactly the case with Star Core. Lucent and Motorola started collaborating on their joint DSP design more than a year ago; the formation of Star Core simply formalizes their efforts. The new DSP family will be incompatible with existing DSPs from either company. Assembly-level source translators will be created, but no binary compatibility will exist among any of the companies' DSP families. Within the new architecture, different core implementations will also exist, with some cores implementing just a subset of the features.

The deal is similar to the one struck between Hitachi and SGS-Thomson (see MPR 12/29/97, p. 10), whereby the two companies will jointly define an instruction-set architecture but separately develop and market the derived devices. Motorola and Lucent executives cited the usual market targets for their new chips: communications, transportation, and consumer electronics.

The two companies have also agreed to cross-license three existing families. Motorola gets access to Lucent's DSP16000 architecture, while Lucent gains a license to the 56800 and M•Core designs (see MPR 10/27/97, p. 12). The Star Core agreement will not alter product roadmaps for any of these families; Motorola's stated intent to add DSP capabilities to M•Core will proceed from the Star Core facility.

The agreement, which has a four-year life span, suggests that more than one new DSP family may be created from the union. After the first architecture is fully defined, and one or two core implementations finished, the group may move on to another new architecture in a short time.

While the agreement is a big deal for both companies, it may not qualify as a "revolutionary announcement in the history of the semiconductor industry," in the words of Hector Ruiz, president of Motorola's chip operations. By joining forces, the #2 and #3 DSP vendors can better head off TI as it moves inexorably toward DSP market dominance. Worldwide unit sales and revenues of DSPs are growing even faster than those of microprocessors; by combining their efforts early, Motorola and Lucent stand a better chance of sharing a strong position going into the next decade. —*J.T.*

Demise of Somerset Splits Embedded Vendors

With the sun setting on Somerset, the once bucolic design center jointly founded (and funded) by IBM and Motorola (see MPR 6/22/98, p. 4), the course of embedded strategy from these two companies will take a turn. IBM has essentially ceded the midrange of the embedded PowerPC market to Motorola, leaving the latter company to focus on communications controllers while IBM pursues ASIC business.

With Somerset, IBM had a three-pronged strategy that included servers, midrange processors (primarily for Apple), and its 40x series of ASIC cores. Without Somerset, IBM will have no midrange processor cores to replace the existing 740 and 750 (G3) series. Instead, IBM's embedded design center near Raleigh (North Carolina) plans to "cherry pick" CPU designs from the company's three high-end design centers in Rochester (NY), East Fishkill (NJ), and Austin (Texas).

Motorola will continue using Somerset to develop new midrange embedded processors, especially new cores for its line of integrated communications and networking chips. While IBM fine-tunes its cores for ASIC development, Motorola will produce packaged parts. Thus, the two companies will split along custom/commercial lines, as well as over microarchitectural implementations.

For the short term, both companies still have access to the 603e/604e and 740/750 processor cores and will likely use them in ASIC designs (IBM) and custom controllers (Motorola). Beyond about 1999, however, the two companies will be working with separate cores. In general, we expect IBM's processors will have somewhat lower performance than Motorola's, which will follow the previously published PowerPC roadmap more closely (see MPR 8/26/96, p. 12). Even with different cores, software compatibility should not suffer after the two companies part company. —*J.T*

VLSI Spins ARM-Based CDMA Controller

Original ARM investor and founding member VLSI Technology is extending its line of telecommunications-related controllers with a new chip for CDMA digital cellular telephones. The announcement comes on the heels of VLSI's recent GSM chip introduction (see MPR 12/8/97, p. 9) and positions the company well for an expected upsurge in sales of digital handsets in North America.

The new chip, dubbed CDMA+ Processor 100, contains an ARM7TDMI core, an Oak DSP core, Bytes of on-chip ROM and RAM, a CDMA/AMPS radio, and various peripheral and memory interfaces. To build a CDMA handset, little more than an RF front end, memory, keypad, and battery are required. The device is housed in a plastic BGA package that measures just 12 mm on a side.

The chip's dual-processor architecture requires two code sets to run. VLSI includes royalty-free Oak DSP code with the device, some of which is contained in the on-chip

ROM. The remainder must be stored off chip. VLSI believes that some of its CDMA firmware is still in flux and is therefore hesitant to commit the entire code set to ROM in the current chip. Future devices may well include more ROM with all the code. VLSI separately licenses the companion code for the ARM core with a one-time fee and no royalties.

The ARM core can run at 16, 20, or 24 MHz, depending on how much performance headroom customers want for user features. The Oak core always runs at 50 MHz. As the CDMA market grows and processing demands increase, future versions of the chip are likely to replace the Oak DSP core with a Palm core for higher performance. The ARM7 core is likely to remain, but with higher clock rates. In time, the ARM9 may appear for high-end handsets, perhaps those with data-communications ability.

The CDMA+ chip is aimed directly at Qualcomm, the company's major competitor in the burgeoning CDMA semiconductor business. VLSI's major competitive tactic is "we're not Qualcomm," meaning that VLSI does not compete with its potential customers for cellular-handset business. VLSI is also in the business of making customer-specific devices, which may appeal to large companies like Nokia, Motorola, or Ericsson.

It's curious that VLSI, which is a major investor in ARM Holdings, did not choose ARM's Piccolo coprocessor for its DSP capability (see MPR 11/18/96, p. 17). The company believes that Piccolo doesn't have enough performance for its current needs or the growth potential to satisfy future wireless standards. For the foreseeable future, VLSI is staying with separate CPU and DSP cores for its wireless offerings.

VLSI's local market for digital cell phones is still relatively small, as North American countries are shifting from analog to digital wireless service more slowly than European and Asian countries. That situation is expected to change rapidly, however—a shift that is particularly hurting Motorola—as CDMA and TDMA telephones become the norm. Market-research reports place the analog/digital crossover point early in 1999, with CDMA penetration at 50% (nearly 50 million units) by 2001. At this rate, VLSI has a large and lucrative market ahead of it. —*J.T.*

PowerPC EC603e Hits 300 MHz

Motorola has decorated the top of its family tree with a 300-MHz version of its EC603e embedded microprocessor (see MPR 10/6/97, p. 8). The company's embedded line now matches the desktop PowerPC 603e clock-for-clock. As the desktop market for PowerPC moves to the PowerPC 740 and 750 chips, the 603e becomes Motorola's platform for high-end embedded systems.

The company now offers this chip in no fewer than seven speed grades, hitting every 33-MHz step from 100 MHz to 300 MHz. IBM has matched only a few of these clock speeds with its identical EM603e chip. All versions of the EC603e (and EM603e) are plug-compatible, except for differences in supply voltage.

The newest speed grade is much more expensive, relatively speaking, than its siblings. Whereas the slower EC603e chips sell for about $0.21 per MHz, the 300-MHz chip's $109 price works out to $0.36 per MHz. Like Intel with its Pentium II, Motorola charges a steep premium for its fastest part, out of proportion to the chip's performance. This price is also not much of a discount from that of the full-featured PowerPC 603e, which sells for $135 at the same clock rate and quantities.

While there aren't many embedded processors that cost $100, there are virtually none that run at 300 MHz. That price puts the EC603e in the same league as the R4700 or R5000, both of which have working floating-point units but 50% slower clock rates with proportionally lower Dhrystone MIPS ratings. It is much cheaper, however, than the ridiculously overpriced 68060 (and 'EC060 and 'LC060) or Intel's 266-MHz Pentium II embedded module (see MPR 6/1/98, p. 15). The PowerPC's cache snooping might give it a slight edge over the MIPS parts, but only if floating-point isn't needed. The EC603e offers much better integer performance for the price, but the MIPS chips are the only choice if floating-point code is used. For designers upgrading from other PowerPC chips, the new EC603e is a competitively priced top end to Motorola's high-end embedded lineup. —*J.T.*

Motorola Scraps Celestri Satellites for Teledesic

Motorola has scrapped plans for its own $13 billion Celestri satellite network and joined forces with the likes of Bill Gates and Craig McCaw in backing Teledesic, a rival data-satellite network. The company is investing $750 million in Teledesic, making it a 26% partner in the private venture. Motorola's stake in the existing Iridium satellite network is unaffected.

The move comes as Motorola prepares to decimate its workforce (laying off some 15,000 employees) and post a nearly $2 billion loss in 1998. Most of the responsibility for the year's disappointing results was laid at the feet of the company's semiconductor operations, which in turn blamed Asian economic turmoil for the flagging sales.

Celestri, like Teledesic, was to have begun launching satellites in 2002 that would stay in low-earth orbit (as opposed to higher geosynchronous or elliptical orbits). Although Celestri was to have consisted of only 64 satellites, it would have been more costly than Teledesic, which plans to launch 288 relatively simple satellites. Both networks were to have carried primarily data communications, with some voice traffic.

About 1,300 commercial satellites are scheduled for launch over the next nine years, an average of two or three new satellites per week, all privately owned. Companies like Loral, Motorola, and Hughes, and private investors like Gates and McCaw, are investing billions of dollars in extending the wireless infrastructure globally. When these competing networks begin functioning in a few years, the market for wireless telephones, data terminals, and other devices will grow even faster than it is now. —*J.T.*

Image Sensors

Continued from page 7

blue, and two green filters is often used, arranged in a Bayer checkerboard pattern (two green filters due to the eye's greater sensitivity to green light). Then, a digital signal processor interpolates the intensity values of adjacent sensors to construct an RGB color estimate for each pixel in the image.

After the image is colorized, the digital signal processor performs additional functions to create the final image. These functions include exposure correction, removal of residual distortions from the lens system, image enhancement, white balancing (to correct for various light temperatures), and gamma correction (to match the linear response of the sensors to the logarithmic response of the eye).

Other functions necessary to implement a complete camera system include autofocus, image sharpening, electronic pan and zoom, image compression (e.g., JPEG), and communication of the image to the camera's memory, an LCD screen, a PC, or a television.

CCDs Provide Best Images

The first digital image sensors on the scene were charge-coupled devices (CCDs). CCDs operate by collecting charge generated by the photoelectric effect in a buried-channel MOS capacitor. A CCD pixel is formed by multiple gates, as Figure 2 shows, that are held at different voltages to create a potential well for collecting charge.

To read data in the sensor array, carefully timed clock signals are applied to the gates to march the charge from one pixel to another down the array—hence the name charge-coupled device. As a row of pixels is shifted out of the array, the charges are dumped into an analog shift register and shifted out serially to be converted to a voltage signal.

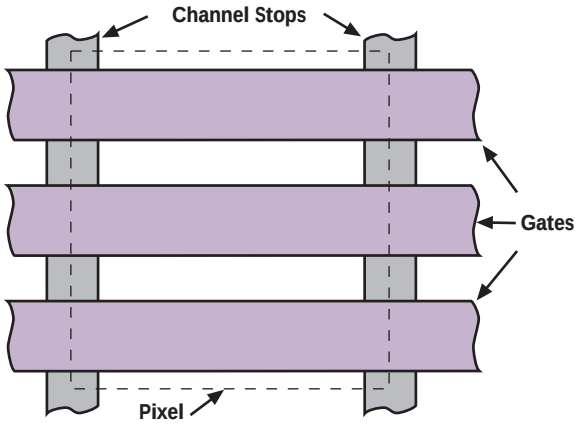

Figure 2. In a CCD, photons striking the pixel are converted to charge by the photoelectric effect and stored in the depletion-region beneath the gates formed by the gate potentials. These charges can be transferred from pixel to pixel down the array by modulating the gate voltages with carefully timed clock pulses. Channel stops isolate pixels in one column from those in the next.

In a CCD, most of the silicon area is devoted to photon collection. The high fill factor (the ratio of photosensitive area to pixel size) gives CCDs good photoelectric efficiency. The FET structure and semiconductor process used for CCDs are designed to make good capacitors, maximize SNR, and achieve high charge-transfer efficiency (up to 99.999%). Achieving these goals requires a specialized process with large voltage swings and multiple supply and bias voltages.

A CCD's high photoelectric efficiency permits pixels to be packed tightly, producing high-resolution arrays on reasonably sized die. Today's megapixel CCDs typically use 4- to 5-micron pixel pitches on 0.5-inch square die (160 mm^2).

While CCDs have good efficiency, they do not scale well with technology. The minimum size of the pixel is dictated by the area required to collect photons and the diffraction limits of the optical system, not by circuit feature sizes. Improvements in process technology can boost photoelectric efficiency, allowing the pixel to shrink somewhat. But these improvements are more gradual than the square-law improvements we expect from process shrinks.

CMOS Image Sensors Use Active Pixels

Unfortunately, the very architecture that makes CCDs efficient also makes them expensive. Their FET structure is incompatible with modern CMOS VLSI processes. As a result, CCDs cannot take advantage of the industry's massive CMOS fabrication infrastructure. The CCD process is complex and thus suffers from poor yields, further increasing costs.

Using pure-CMOS to overcome the shortcomings of CCDs is not a new idea. But a narrow dynamic range and high noise levels have plagued CMOS sensors until recently.

Early CMOS sensors used a passive-pixel structure (PPS). In this arrangement, a photodiode is paired with a transistor switch, as Figure 3 shows, to dump the accumulated charge onto a column bus. This structure has advantages over CCDs in that the array can be X-Y addressed and read out a row at a time. This method lowers the analog bandwidth requirement compared with the serial readout of CCDs, reducing susceptibility to noise injection. But the extra transistor takes up space, which lowers the fill factor and gives PPS sensors lower efficiency and SNR than CCDs.

Recently, however, CMOS scaling has made transistors small enough that an amplifier can be included at every pixel. Although the amplifier takes area and reduces the fill factor even further, the amplification more than compensates. This active-pixel-sensor (APS) technology raises the SNR and efficiency of CMOS image sensors near those of CCDs, bringing image quality above the threshold necessary to serve the consumer digital-camera market.

Although CMOS photodetectors do not scale any better than CCD photodetectors, CMOS pixels will scale better than CCD pixels. As CMOS transistors shrink, additional active circuitry can be placed at each pixel sensor. With 0.18- and 0.13-micron technology, each pixel can include multistage amplifiers, ADCs, and even digital signal processing.

Such circuitry can increase sensitivity and reduce noise, allowing the pixel to shrink. The overall pixel will still not shrink at the rate of normal CMOS circuitry, but it will scale better than its CCD counterpart.

CMOS Allows Integration

With image-quality problems licked, the integration, power, and cost advantages of CMOS come into play. With CMOS, unlike CCDs, the ADCs can be integrated onto the die, close to the sensor array. This design reduces the analog signal's susceptibility to noise and makes the output of the chip fully digital. CMOS also makes it possible to integrate on a chip digital signal processing and other functions necessary to produce a complete digital camera.

CMOS sensors require less power than CCDs. They operate on a single low-voltage supply. (On-chip bias generators can provide different voltages for the sensors or ADCs if necessary.) All clock signals can be generated and contained on chip, and all off-chip communication is fully digital. The net effect is dramatically lower system power requirements than CCDs—a big plus for battery-powered cameras.

CMOS sensors also beat CCDs in manufacturing costs. Standard CMOS processing allows CMOS sensors to take advantage of the enormous infrastructure and learning curve of the semiconductor industry. CMOS sensor processes are simpler than CCD processes, giving them higher yields. This yield becomes especially important as the arrays are scaled to higher resolutions. CMOS sensors are also less susceptible to defects. A defect may affect only a single pixel in a CMOS sensor, whereas in a CCD, a similar defect would wipe out an entire column. These factors make CMOS image sensors more manufacturable than CCDs.

CMOS sensors should ultimately have a price advantage over CCDs as well. Today, there are about five large manufacturers of CCDs. With CMOS sensors, anyone with a 0.35-micron fab can—and probably will—get into the market. The competition from many vendors should drive prices down.

Recently, prices of CCDs have fallen dramatically, driven down by the competitive threat from CMOS sensors. Low-end 320 × 240 CCDs are now available in the $12 to $35 range, and 640 × 480 versions have dropped below $80 (plus another $10 for external electronics). CMOS sensors are already slightly less costly than CCDs at these resolutions.

High-end 1,000 × 1,000 CCDs capable of 30 fps are still pricey at over $500, but these will be coming down as manufacturers struggle to squeeze every last drop of return out of their CCD investments. Because of their better yields, CMOS sensors are expected to be far less expensive than CCDs at these higher resolutions. Eventually, the manufacturing cost advantages of CMOS will become overwhelming, and CCDs will be unable to keep up the pace.

Several Companies Making CMOS Sensors

Today, there are relatively few CMOS sensors on the market, but there is an enormous amount of activity, and we expect a continuous stream of new announcements. Unfortunately, details of many of the upcoming sensors are unknown at this time. Table 1 presents a representative sampling of recent activity.

Intel sells the 971 PC Camera Kit, which contains its own 768 × 576-pixel CMOS APS sensor. The 80971AC is built in the company's standard 0.35-micron P854 process. Sources in the imaging industry believe that Intel's sensor technology was developed by Photobit (La Crescenta, Calif.). Although neither company acknowledges this relationship, it seems likely that Intel would try to leverage existing technology rather than develop new technology itself.

Intel has announced a major initiative with Kodak to jointly develop imaging products. The agreement includes a broad patent cross-license. The companies agreed to spend $150 million over three years to promote their imaging products. Sources indicate that image sensors are not part of the Intel-Kodak agreement.

Kodak has, however, teamed up with Motorola to build image sensors based on technology it calls ImageMOS. Photobit is the source of that technology and is also the design house for Kodak-Motorola sensors. Kodak says it has several cameras using this technology on the drawing boards.

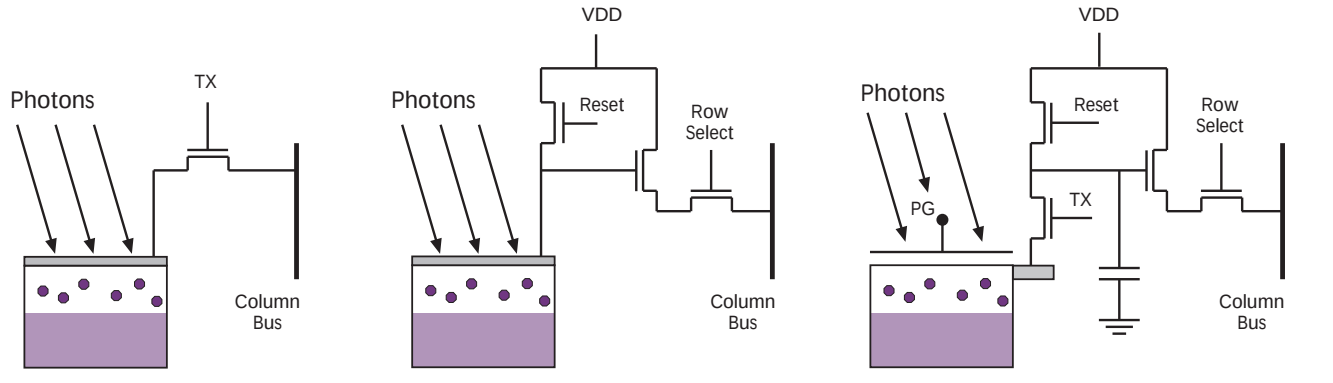

Figure 3. Early CMOS image sensors used a passive-pixel arrangement. But CMOS scaling now allows active pixels, which have amplifiers at each pixel site to improve the light sensitivity and signal-to-noise ratio.

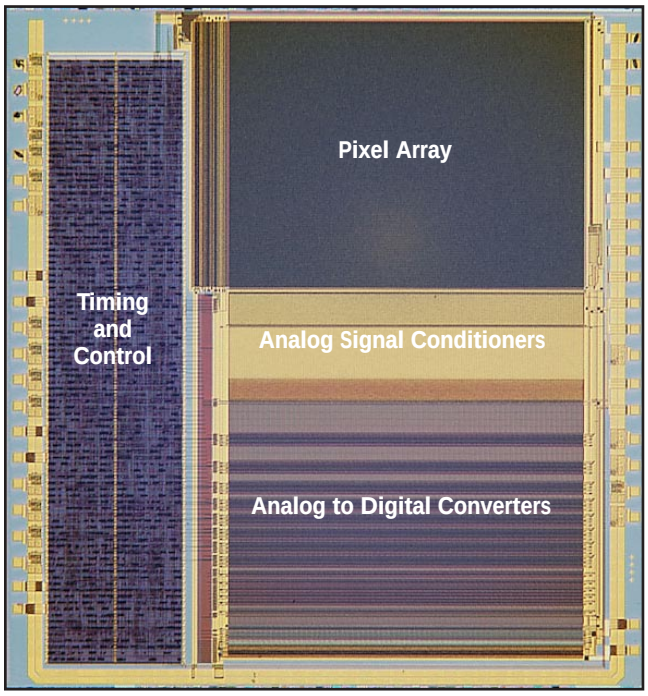

Figure 4. Photobit's PB159 CMOS image sensor is 8.6 × 9.4 mm with 196,608 (512 × 384) active pixels and column-parallel analog-to-digital converters. Each pixel is 7.9 × 7.9 microns.

Photobit was spun out of NASA's Jet Propulsion Lab (JPL) in 1995 with an exclusive license to JPL's APS technology from Caltech. At Hot Chips last year, Photobit described its APS technology in conjunction with a 512 × 384-pixel APS. Figure 4 shows the PB159, which uses a pixel consisting of a photodiode and three transistors similar to the design in Figure 3(b). The PB159 integrates 8-bit column-parallel ADCs, automatic exposure logic, windowing, and an I^2C serial command interface. Presumably, it is similar to the technology employed in the Kodak, Motorola, and Intel sensors. Photobit also has a 1,280 × 1,024 CMOS APS with an on-chip 10-bit ADC that is capable of 60 fps.

Rockwell Semiconductor has introduced five devices ranging from 352 × 288 pixels (CIF resolution) to 960 × 720 pixels. The Ri0352A, 640A, 800A, and 960A all have APS, while the Ri0352P does not. The sensors all integrate signal-conditioning circuits, a 10-bit ADC, and a digital video interface capable of 30 fps. There are rumors that the Rockwell technology may also be based on Photobit's. Prices for the Rockwell sensors range from $18 to $63.

A unique feature of the Rockwell sensors is "microlenses." These are 0.5- to 0.8-micron glass beads that concentrate light onto the detectors and raise the SNR to 46 dB. We expect microlenses to be a feature of many of the new sensors coming from other manufacturers.

VLSI Vision (Edinburgh, Scotland) supplies a 1,000 × 800-pixel CMOS APS to Sound Vision (Framingham, Mass.) for its SVmini-2 camera, which is marketed by Vivitar and Umax. Both of these cameras sell for around $400. At ISSCC98, the company described a more interesting 306 × 244-pixel single-chip NTSC video camera that includes a 300-MOPS dataflow processor to perform signal processing. The 68-mm^2 chip uses a 12 × 11-micron pixel and is implemented in 0.8-micron double-layer-metal CMOS. VLSI Vision said that with microlenses, low-power design techniques, process improvements, and a process shrink, it plans to get the SNR up to 57 dB (from 45 dB currently).

Texas Instruments announced it will sample the TC286 640 × 480-pixel image sensor in July. TI, a major supplier of CCDs, claims its CMOS sensor has image quality comparable to that of CCDs. In future chips, TI plans to integrate ADCs, digital signal processing, and image compression.

Toshiba has demonstrated a prototype 1.3-million-pixel CMOS APS that meets the requirement of the 1,280 × 1,024 SXGA format. The die measures 144 mm^2 in 0.6-micron triple-poly double-level-metal CMOS and uses a pixel size of 5.6 × 5.6 microns. The triple-poly process is more complex than what other manufacturers are using but may be responsible for the device's small pixel size. At ISSCC this year, the company said it had further reduced the pixel size to just 3.7 × 3.7 microns in the same process.

Other Designs in Progress

AT&T Bell Labs (now Lucent Technologies) also entered into a technology agreement with Photobit. The company recently described a 354 × 292-pixel single-chip digital camera based on a five-transistor photogate APS shown in Figure 3(c). Color reconstruction and gamma correction are performed with 100 MOPS of on-chip digital signal processing.

Atmel and Polaroid have joined forces to develop CMOS image sensors that will integrate analog signal-conditioning

CMOS APS Device	Process µm	Process Poly/Met	Resolution	Pixel µm	SNR dB	ADC bits	Integration	Rate fps	Power mW	Power Volts	Comments
VLSI Vision	0.8	1P2M	306 x 244	12 x 11	45	10	300 MOPS DSP	30	550	5	Camera on a chip
Bell Labs, Lucent	0.8	1P2M	354 x 292	18 x 18	n/a	8	100 MOPS DSP	30	182	3.3	Shown at ISSCC98
Photobit PB159	0.5	2P3M	512 x 384	7.9 x 7.9	60	8	Autoexposure	30	50	5	Column-parallel ADCs
Intel	0.35	854	768 x 576	n/a	n/a	n/a	n/a	still	n/a	n/a	In 971 PC Camera Kit
Rockwell Ri0960A	n/a	n/a	960 x 720	n/a	46	10	Serial I/F	30	100	3.3	Uses microlenses
VLSI Vision	0.8	1P2M	1000 x 800	10.5 x 10.5	64	12	None	5	100	5	SVmini-2 camera
Photobit PB720	0.5	2P3M	1280 x 720	7.9 x 7.9	60	10	Windowing	60	250	5	
Toshiba	0.6	3P2M	1318 x 1030	5.6 x 5.6*	n/a	n/a	n/a	still	30	3.3	* 3.7 x 3.7 at ISSCC98

Table 1. Nearly all recent CMOS image sensors use active-pixel sensors. New development activity seems to be focused on higher resolution, better SNR (image quality), and higher levels of integration.

circuits, DSPs, and flash memory. Hyundai Electronics has developed an XGA resolution (1024 × 768) CMOS sensor with integrated 6-bit automatic gain control and 8.5-bit ADC.

New Ideas Abound

G-Link (Santa Clara, Calif.), through a partnership with the Institute for Microelectronics Stuttgart (IMS), is developing CMOS sensors based on that company's high-dynamic-range CMOS (HDRC) technology. HDRC delivers 120 dB of dynamic range using a logarithmic pixel-compression technology. This technology eliminates the need for shutters, aperture control, automatic gain control, and gamma correction. G-Link has wafer supply contracts with Chartered Semiconductor and TSMC. It also plans to build SVGA sensors in IBM's 0.35-micron CMOS process.

Amain Electronics (Simi Valley, Calif.), a developer of specialized infrared image sensors, has a unique continuous-modulation scheme for reading data from the array. Instead of dumping all the collected charge at once, as most APS devices do, the charge in the well is continuously sampled. Between sampling operations, a fixed amount of charge is switched into a subtraction capacitor. Amain says this approach, along with its oversampled sigma-delta ADC, has a smoothing effect that reduces noise.

Stanford University's Image Sensor Group has demonstrated pixel-level ADC. In this approach, an analog-to-digital converter is placed at every pixel. Each ADC uses very little power, because it has to run at only a few kilohertz. Since all analog signals are confined to the vicinity of the pixel, noise is reduced. Another advantage of this approach is scalability; the same pixel can be used for arrays of any size. The ADC requires more area than an APS amplifier, but Stanford has demonstrated a 9 × 9-micron pixel in 0.35-micron technology with a good (25%) fill factor. Both HP and Intel have built test chips.

The IMEC research center in Leuven, Belgium, is developing a technology that uses special p+ doping to create an electrostatic barrier around the pixel. The barrier funnels carriers generated anywhere in the pixel area into the junction that is accumulating charge. This technique should raise the sensitivity of CMOS sensors close to that of CCDs.

CCDs to Fade Away

CMOS image sensors will not eliminate CCDs soon. CCDs still produce higher-quality and higher-resolution images, and that will likely remain true until support for CCD research finally dries up. The industry has 25 years of experience manufacturing and using CCDs, whereas CMOS sensors have become viable only in the past three or four years. CCD development continues: Sony, Matsushita, NEC, and TI have all recently announced new megapixel CCD sensors.

But the CMOS sensors coming onto the market now are every bit as good as CCDs at the low end. Nearly all the CMOS sensor manufacturers have on the drawing boards megapixel versions that will rival the best CCDs in resolution and image quality within the next two to three years.

For More Information

For more information on CCDs or CMOS image sensors, access the following Web sites:

- Intel, *www.intel.com/imaging*
- Photobit, *www.photobit.com*
- Motorola, *mot2.indirect.com/adc/markets/image.html*
- Rockwell, *www.nb.rockwell.com/platforms/personal_imaging/*
- TI, *www.ti.com/sc/docs/disp/disp.htm*
- G-Link, *www.glinktech.com*
- VLSI Vision, *www.vvl.co.uk/*
- Stanford Image Sensor Group, *www-isl.stanford.edu/people/dyang/imaging.shtml*
- Sony, *www.sel.sony.com/semi/ccdarea.html*

The advantages of CMOS image sensors over CCDs seem compelling. Standard CMOS processes have lower wafer cost and better yields. As competition heats up—which is bound to happen, since CMOS sensors can be built by just about anyone—prices may drop even faster than costs. Integration of other circuits on the same die as the sensor array has obvious advantages in power, space, and system-level costs. By Photobit's estimates, CMOS APS sensors need 1% of the system power and 10% of the physical volume of CCDs with comparable image quality.

As CMOS geometries shrink, integration of digital processing horsepower, at each pixel or next to the sensor array, will offer capabilities beyond that possible with CCDs. Near-35-mm-film resolution and MPEG2-encoded HDTV video should be available in a few years from a single chip in your PDA. On-chip DSPs performing billions of operations per second will eliminate exposure, focus, stability, and lens-distortion problems, making good photographers out of rank amateurs.

Intel sees digital imaging as a MIPS-intensive application that can showcase its high-performance microprocessors and drive demand for PCs. But beyond that, its establishment of a full division with profit-and-loss responsibility indicates that it also sees digital imaging as a big business in and of itself. Intel wants to be a vertical supplier to this market, providing sensors, microprocessors (e.g., StrongARM), ASICs, and flash memory. CMOS sensors provide a vehicle for putting aging fabs to good use, and Intel is betting heavily on them in its bid for this market. Intel's presence alone should make CCD manufacturers nervous.

Today, CCDs still deliver superior image quality over CMOS image sensors. This will remain true for many years in applications where power, space, and cost are not issues. So if you need a cryogenically cooled infrared sensor for your space telescope, look to CCDs—otherwise, your best bet will be a CMOS image sensor. M

Testing Reveals x86 Core Differences

Design Tradeoffs Can Create Anomalies

by Brian Case

At the same clock rate, all the leading x86 cores—Pentium, Pentium Pro, Pentium II, and K6—have similar performance on most desktop applications. Application-level benchmarks like Winstone and synthetic benchmarks like WinBench typically do not reveal large differences among these cores.

But the broad scope of these benchmarks masks more significant differences underneath. Peeling the onion back one layer reveals that these cores have performance differences as large as 2× on specific CPU-intensive tasks. Even more surprising, the direction of the difference can be counterintuitive. We found, for example, cases where the venerable Pentium core takes fewer clocks than its more complex out-of-order cousin, Pentium II.

For those interested only in average performance on PC productivity applications, the traditional benchmark results provide reasonable guidelines. But for those looking to use a CPU for a specific task, and those interested in microarchitectural tradeoffs, it may be worthwhile to look deeper.

Tuning for One Application Domain Can Backfire

An effort to harness the computing power inherent in the servers and clients that form the Internet is being coordinated at *www.distributed.net*. The effort has discovered, through brute force, the secret keys to a couple of cryptography challenges. Available at the Web site are client key-checking programs that coordinate to check every possible key in the key space until the one correct key is found. The clients implement key-checking for the DES and RC5 encryption algorithms.

These client key-checkers are available for many types of processors, and the x86 clients allow the user to select core code that has been hand-tuned for the most popular x86 chips. While it is beyond the scope of this article to verify that a tuned version is optimal for a given x86 chip, a quick check indicates that the tuned versions do indeed yield good performance on the intended chip. Table 1 reports the results of running the respective hand-tuned versions.

Testing the performance of these key-checking clients on various processors gives surprising results. Table 1 shows that when compared at equal clock rates, the fastest DES key-checker is the Pentium/MMX. The complexity of the P6 core pays better dividends on the RC5 benchmark, running essentially twice as fast as the Pentium. Pentium Pro and Pentium II (Klamath, in this case) have essentially identical performance.

Tests with the L2 caches disabled and various tunings of the BIOS memory timing parameters produced identical results. This indicates the benchmark fits in the L1 caches, so the performance of these benchmarks depends only on the microprocessor core and L1 caches.

While it may be startling to see the P55C microarchitecture beating the dramatically more sophisticated P6 core, the real surprise is K6, which performs poorly on these tests.

Investigation indicates that the K6 performance problem on this test stems from a single design factor: the implementation of the rotate instructions. Since these encryption algorithms use a lot of shifting and rotating of long bit strings, a slow implementation of rotates can significantly reduce performance. The K6 implements the rotate-register instruction with a microcoded ROP sequence, while the P6 implements it with a single internal ROP.

Equalizing Clock Rate Can Be Misleading

The results shown in Table 1 are both revealing and fallacious. The test is revealing if the goal is to expose the relative efficiencies of different microarchitectures.

The fallacy arises because the P6 microarchitecture was designed with a very long pipeline to enable operation at higher clock rates, resulting in what appear to be inefficiencies relative to a shorter pipeline. Note that the same fabrication process that yields 233-MHz Pentiums produces 300-MHz Pentium II CPUs, and a A 300-MHz Pentium II will soundly trounce a 233-MHz Pentium/MMX.

Pentium Extensions Ease Performance Monitoring

With the Pentium processor, Intel introduced a host of performance monitoring and debugging capabilities. Intel added even more in the P6. Most measure specific events such as L1 cache misses or mispredicted branches.

While Intel's competitors have not copied all these capabilities, one very useful feature seems ubiquitous: the Time Stamp Counter (TSC). The TSC value advances at the processor core clock rate and is always running. The RDTSC

Microprocessor (all @ 233/66 MHz)	DES-II		RC5-64	
	Mkeys/s	% Max	Mkeys/s	% Max
Pentium II	1.09	91%	**0.656**	**100%**
Pentium Pro*	1.08	90%	**0.654**	**100%**
K6	0.82	68%	0.381	58%
Pentium/MMX	**1.20**	**100%**	0.334	51%

Table 1. The DES-II and RC5 decryption clients are CPU-bound programs that fit in the L1 caches; thus L2 cache, memory, I/O, and display performance are irrelevant. *Pentium Pro-200 overclocked to 233 MHz. (Tests run as "rc5des -c n -benchmark" using client version 2.7020.403 under Linux.) (Source: MDR)

instruction reads the value of the 64-bit TSC into EDX:EAX. By sampling the TSC with a pair of RDTSC instructions, a program can measure the elapsed number of cycles between two points in a program (see sidebar, next page).

Characterizing Rotate Performance

Table 2 shows the measurements obtained using RDTSC to count the execution times of several loops.

The P6 core executes the empty decrement/jump-not-zero loop slower than the other processors. With branch prediction and the capability to retire two or more instructions per cycle, the natural assumption is that a two-instruction loop will execute at the rate of one iteration per cycle, as shown by the K6 and Pentium/MMX, but the P6 can decode branches only in its first decoder position. In real programs, empty loops are not very useful, so the suboptimal performance of the P6 in this case is unimportant.

As the table shows, the P6 can overlap up to two independent ROL instructions with the loop overhead. Although not shown, a third ROL adds a single cycle because the P6 has only one shift execution unit. The inclusion of two independent increment instructions adds only a single cycle to the count, since the P6 has two execution units capable of integer addition; P6 performance is unaffected by the static scheduling of the instruction sequence (columns 4 and 5).

Each ROL adds a cycle to the count for Pentium/MMX, because ROL is not a pairable instruction. Even though Pentium/MMX has two integer execution units, the static scheduling of the sequence with two ROLs and two INCs matters, since the in-order microengine has no capability to reorder instructions.

Each ROL adds two cycles to the total for the K6, because ROL is a microcoded instruction. Not only does the microcoded sequence of ROPs take longer, but it also monopolizes the decoders while the ROP sequence is generated, as illustrated by the difference in the cycle counts in columns 4 and 5. In the first sequence, the first ROL stalls the decode, but the two INCs get decoded together and executed in parallel. In the second sequence, with the ROL and INC instructions interleaved, all four ROL and INC instructions are decoded separately. The ROLs and INCs are independent, but this decoder bottleneck starves the execution units.

The last two columns in Table 2 show the effect of adding one and two more independent INCs to the sequence from column 4. For the K6 and Pentium/MMX, the additional INCs in columns 6 and 7 do not add to the execution time, because they are paired and can use otherwise idle decode and execution resources.

The oddity in the cycle count for the P6 core on the code in column 6 indicates the sequence executes in three cycles for half the iterations and in four cycles on the other half. This is probably caused by instruction pairing across the loop boundaries.

The IPC (instructions per cycle) columns in Table 2 are a measure of how effectively the execution hardware is used. For these code sequences, only the K6 and Pentium/MMX ever achieve their theoretical maximum of two instructions decoded, executed, and retired per cycle. The P6 core has a theoretical maximum of three IPC, but this peak is not achieved for these sequences. With two integer execution units and one load/store unit, P6 can only achieve three IPC when loads or stores are present. The P6 does, however, exhibit the most consistent performance, which illustrates the power of a general out-of-order microarchitecture.

Winstone results indicate that the K6's poor IPC numbers on these microbenchmarks are not representative of the chip's performance on common PC applications; rather, they are a direct result of a design tradeoff that sacrifices the performance of rotate instructions. For the vast majority of

Test Loop Instruction Sequence	1 L1: DEC EAX JNZ L1		2 L1: ROL EBX,3 DEC EAX JNZ L1		3 L1: ROL EBX,3 ROL ECX,3 DEC EAX JNZ L1		4 L1: ROL EBX,3 INC EDI INC ESI ROL ECX,3 DEC EAX JNZ L1		5 L1: ROL EBX,3 INC EDI ROL ECX,3 INC ESI DEC EAX JNZ L1		6 L1: ROL EBX,3 INC EDI INC ESI ROL ECX,3 INC EDX DEC EAX JNZ L1		7 L1: ROL EBX,3 INC EDI INC ESI ROL ECX,3 INC EDX INC EBP DEC EAX JNZ L1	
Processor	**Cycles**	**IPC**	**Cycles**	**IPC**	**Cycles**	**IPC**	**Cycles**	**IPC**	**Cycles**	**IPC**	**Cycles**	**IPC**	**Cycles**	**IPC**
Pentium Pro	2	1	2	1.5	2	2	3	2	3	2	3.5	2	4	2
Pentium II	2	1	2	1.5	2	2	3	2	3	2	3.5	2	4	2
Pentium/MMX	1	2	2	1.5	3	1.33	4	1.5	5	1.2	5	1.4	5	1.6
K6	1	2	3	1	5	0.8	6	1	7	0.86	7	1	7	1.14

Table 2. The empty loop executes faster on the K6 and Pentium/MMX than on the P6-based processors. The increasing cycle counts for the K6 show how a microcoded instruction reduces efficiency by stalling decode. The fourth and fifth code sequences illustrate the advantage of out-of-order execution: the P6 can reorder the instruction sequence to achieve maximum performance, whereas the Pentium/MMX cannot. (Source: MDR)

Cycle Counting With RDTSC

The Time Stamp Counter and the RDTSC instruction were added to the x86 architecture with the introduction of the Pentium processor. The TSC is a 64-bit counter that is initialized to zero following hardware reset. It then increments by one with each internal processor cycle even when the processor is halted by the HLT instruction or the external STPCLK pin.

Architecturally, RDTSC is guaranteed to return a monotonically increasing unique value on each execution, except in the case when the counter wraps around, which will not occur for years after processor reset.

On an in-order processor, the RDTSC instruction could be used to obtain very accurate measurements of instruction sequences. The architectural definition of RDTSC, however, specifies that RDTSC is not serializing and can be executed out of order with respect to other instructions. Thus, a given RDTSC may be executed before preceding instructions or after following instructions. Theoretically, the uncertainty of the execution order of a RDTSC is determined by the size of the window an out-of-order processor uses to buffer waiting instructions. In practical situations, execution of a RDTSC is fairly predictable due to dependencies.

To filter out uncertainties in the execution order of a pair of RDTSC instructions and to stabilize cache behavior, instruction sequences investigated in this article were incorporated into the following framework:

```
            RDTSC
            PUSH    EAX
            MOV     EAX,1000

            .align  32
L1:         Instruction 1
            Instruction 2
            ...
            Instruction n
            DEC     EAX
            JNZ     L1

            RDTSC
            POP     EBX
            SUB     EAX,EBX
```

An instruction sequence to be measured is inserted in a loop that executes for 1,000 iterations. The RDTSC instructions precede and follow the loop. This strategy has the property that each cycle attributable to an instruction in the loop will cause an increment of 1,000 in the cycle count; any transient cycles—due to the uncertainties in the execution scheduling of the RDTSCs and the overhead of saving the TSC count (PUSH EAX) and setting up the loop count—will appear in the low three digits of the measured cycle count. No uncertainty or transient will amount to more than a few tens of cycles at most.

The execution time for the empty loop (consisting of only DEC/JNZ) is one or two cycles on all processors tested, as Table 2 shows, so loop overhead can be factored out of the final cycle count.

popular PC applications, this tradeoff causes little measurable performance degradation.

Different Caches Perform Differently

Especially for the x86, a fast instruction-execution engine is of little use without a fast data cache to serve the high density of load and store operations in x86 programs. To test the performance of the three levels of a PC memory hierarchy, we designed a tiny loop and a linked-list data structure to allow an arbitrarily long sequence of memory references with controlled cachability.

Figure 1 shows the code sequence, which consists of a load-indirect MOV instruction surrounded by our familiar DEC/JNZ loop. The MOV loads a value from a base register directly into that same base register, so the new value will be used as the memory address on the next iteration. This causes the MOV to be dependent on itself, which ensures that the cycle count for the loop will include the entire execution time of the MOV. It can at most be overlapped with the DEC/JNZ, which has known execution time.

Setting the values of a sequence of memory locations to form a ring, causes this simple code sequence to follow the chain of pointers around in a circle until the count in EAX is exhausted. With the correct number of nodes and addresses of the nodes in the linked list, a cache of a given size and associativity can be overwhelmed, thus causing it to miss on each execution of the MOV instruction.

Table 3 shows the results of using this technique to determine the speed of x86 caches. All processors were running at 233 MHz, and the bus speed of the K6 and Pentium

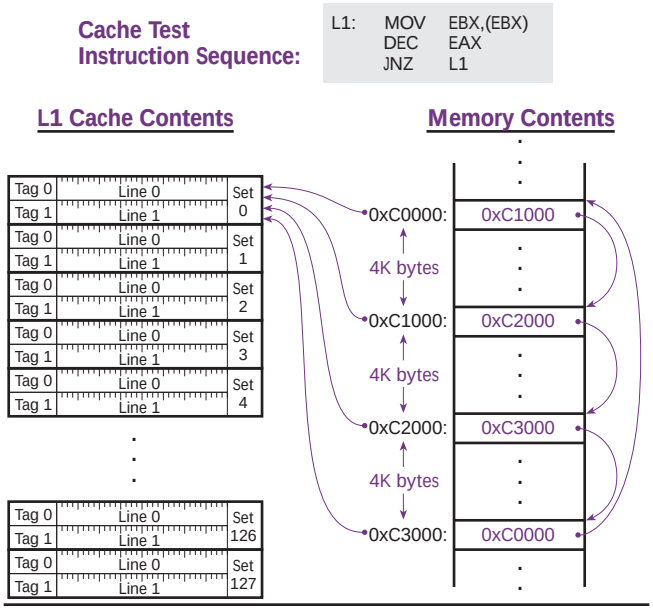

Figure 1. The linked-list data structure is set up such that when the cache test loop runs, the number of elements in the list will exceed the associativity of the cache in which misses are desired. Also, the addresses of the elements in the linked list are chosen so that they map into the same cache set.

MMX was 66 MHz. The cycle counts in the table are for an actual instruction sequence and reflect the true cost of executing worst-case code that delivers an operand to a waiting processor register, not just an isolated parameter such as SRAM access time. In real applications, similar pathological access patterns do arise, but they would not dominate the run time of most nontrivial programs. Database programs, however, may be a significant exception.

The long latencies shown for L2 misses result in a full cache line being loaded, and so, in real programs, the cycle counts in Table 3 are typically amortized over at least a series of sequential accesses to the same cache line.

There are several observations of interest. First, note that the K6 and Pentium/MMX have the fastest L1 caches for 32-bit aligned loads. Any misalignment adds only a single cycle to the execution time for the K6, but Pentium/MMX pays a three-cycle penalty.

The execution time for various alignments highlights one of the improvements Intel made in Pentium II compared with Pentium Pro. Pentium II tolerates misalignment on all accesses except those that cross a 32-byte cache line, but Pentium Pro takes a five-cycle hit when crossing an 8- or 16-byte boundary within a cache line. This Pentium II refinement was probably made possible because its L1 caches are twice the size of Pentium Pro's: the extra 8K of cache can be accessed in parallel, and a multiplexer/shifter can then select the proper group of four bytes. Both processors pay a nine-cycle penalty for accesses that cross a cache line boundary.

The penalty for an L1-cache miss varies greatly, depending on how the L2 is implemented and connected to the processor. The best by far is Pentium Pro, which can deliver a 32-bit operand with less than half the latency of the Pentium II in most cases. Both Pentium II and Pentium Pro processor chips have a dedicated L2 cache bus, which provides fast access to the L2, but Pentium II uses commodity L2 SRAMs, which run at half the speed of the custom Pentium Pro L2 SRAMs. Because its L2 cache is so fast, L1 cache hits for some alignments in Pentium Pro are slower than L1 misses for other alignments.

Once again, examining these processors at equal clock rates may lead to a distortion. The advantage of the Pentium Pro L2 over the Pentium II L2 is probably slightly less than shown in Table 3, because the initial latency of the Pentium Pro custom cache chips is tuned for 200-MHz operation, while the initial latency of the Pentium II cache chips is tuned for 266-MHz (or even 300-MHz) operation. Thus, while the forthcoming Pentium II Xeon CPUs will have a much faster L2 cache than that of the Pentium II, the advantage may be slightly less than Table 3 indicates. For example, an aligned L1 miss/L2 hit may take eight or nine cycles instead of Pentium Pro's seven.

Beware of Benchmarks

Benchmarks like Winstone and Winbench are reasonable indicators of the relative performance of different CPUs and how well they will perform overall. But such benchmarks are designed to test all aspects of PC performance, which makes them poor indicators of how well a processor will perform on a specific task.

For example, as we showed in this article, the K6 may not be the best choice for cryptography applications that depend heavily on rotate operations. On the other hand, the K6 could perform significantly better than the others on applications where misaligned data is prevalent, such as applications that access legacy data structures in network transmission packets.

The results shown here compare cores at the same clock rate. In practice, however, these cores do not all run at the same maximum clock rate. So while Pentium may have better per-clock performance than Pentium II in some cases, Pentium II's deeper pipeline allows it to run significantly faster. This, in fact, may be the more significant contributor to Pentium II's bottom-line performance advantage than its out-of-order microarchitecture.

If you have a specific job in mind for a processor, you may not want to rely on traditional benchmark results to make your decision. By performing a clock-accurate analysis of the processor running code that is representative of your application, you may be able to select a processor that gives you better performance or saves you money. M

Interpretation	Address Alignment (bytes)	K6		Pentium II		Pentium Pro		P55C
		L1 Hit	L1 Miss L2 Hit	L1 Hit	L1 Miss L2 Hit	L1 Hit	L1 Miss L2 Hit	L1 Hit
Naturally aligned on four-byte boundary	0, 4, 8, 12, 16, 20, 24, 28	2	28	3	16	3	7	2
Misaligned, but does not cross 8-byte or 16-byte boundary	1, 2, 3 9, 10, 11 17, 18, 19 25, 26, 27	3	28	3	16	3	7	5
Crosses 8-byte but not 16-byte boundary	5, 6, 7 21, 22, 23	3	28	3	24	8	12	5
Crosses 16-byte boundary	13, 14, 15	3	28	3	28	8	12	5
Crosses 32-byte cache-line boundary	29, 30, 31	3	56	12	54	12	34	5

Table 3. Cycle counts for the 32-bit memory load instruction in the loop shown in Figure 1. All processors are running at 233 MHz; the K6 used a 66-MHz external bus. In some cases, cache performance varies greatly depending on the alignment of the load address. The half-speed cache of the Pentium II more than doubles the L1 miss penalty compared with Pentium Pro. Pentium II Xeon is likely to enjoy a similar advantage over Pentium II. The P55C (Pentium/MMX) results for L1 miss/L2 hit were not reproducible and so are not reported. (Source: MDR)

■ VIEWPOINT

A Swing of the Pendulum?

Future Architecture Changes Unlikely to Improve Performance

by Maurice Wilkes

For some time, there have been mounting signs of uncertainty in the microprocessor world. New forces appear to be at work, and subtle changes are occurring in the climate of opinion. This is true across the range, including processors and their architecture, benchmarks for both processors and software, and simulation. We may have reached one of those periods when we need to take cherished ideas out of our heads, and either discard them or at least dust them off.

At the end of the 1980s, effective benchmarks were developed, and it suddenly became possible to make comparative measurements of processor performance. The result was to focus attention on the processor rather than on the system. Up to that time, computers had been sold on system performance and on the quality of vendor support. No one knew, with any precision, how different vendors' processors compared, nor did they care. Not that customers took no interest in instruction sets; they did, but they judged them like pet dogs for their appearance and pretty points, not like racehorses for their speed. Benchmarking came in the wake of a great boom in simulation. No longer was it necessary to build a processor in order to evaluate an architecture. Given a room full of the fastest machines available, the architecture could be simulated. A great triumph of simulation was to establish beyond any doubt that RISC processors were faster than existing processors by a factor of approximately two.

However, what really gave RISC its chance was not its high speed, but its small size. A RISC processor needed only half as much silicon real estate as a conventional processor. Real estate was still in very short supply, although the situation was improving steadily. By 1989, room could be found on a single chip for a RISC integer processor, an MMU, a TLB, and cache control circuits. The MIPS R2000/3000 processor, designed on these principles, turned out to be fast enough to outperform the fastest minicomputer that DEC then made. Without RISC, this could not have happened for another two years—a conventional processor would by itself have taken up the entire available area, leaving no room for the MMU and other items. These two years were crucial in enabling the Unix workstation based on a RISC processor to establish itself in the market place.

The Unthinkable Happens

It was not long before people stopped buying VAXs and other minicomputers, and began to buy RISC workstations instead. A short time before, most people would have said that the VAX architecture was so firmly entrenched that its demise was unthinkable. But the unthinkable happened, and with so little fuss that the event was hardly noticed. The IBM 370 architecture was similarly hit, although it survived in mainframes.

Silicon real estate is still very expensive—according to Gordon Moore, it sells at a billion dollars per acre. You can still only buy 50 nano-acre lots, but the number of transistors that you can put on a single lot is going up and up. No-one worries any more about real estate as far as the integer processor is concerned.

The minicomputers that the R3000 challenged were based on bipolar circuitry. Thus was another article of faith shattered that, while CMOS was one for personal computers, "real" computers would always be bipolar. Because of this fundamental difference in underlying technology, the personal computer side of the industry had developed separately from the main industry. Now that the difference no longer existed, it was inevitable that PC manufacturers should seek a share of the workstation market. The 486 brought this within sight. The Pentium Pro came very near to the speed of the best RISC chips. This was a notable success and, in order to achieve it, Intel pulled out all the stops known to their chip designers and process engineers.

It is possible that Intel was not confident of being able to repeat this success for smaller feature sizes. Perhaps this is why they have teamed up with HP in the Merced initiative. On the other hand, it may have been a decision based primarily on marketing considerations. At all events, they evidently feel that a design combining RISC and x86 features is what is required. They have chosen to develop an entirely new design, instead of starting, as they might have done, with a proven RISC design, such as the Alpha. They have stated that Merced will be binary compatible with x86. The performance that can be delivered at this level will clearly play an important role in determining the future direction of events.

Future Innovation

In the recent past, increases in processor speed have come partly as a result of architectural improvements and partly as a result of shrinkage. I would be surprised if we were to see any further architectural improvements of a major kind; indeed, most of those that have recently found their way into silicon were foreshadowed, if not fully exploited, in the large mainframes of the 1970s. In my view, any further significant increases in the speed of uniprocessors will come from shrinkage alone.

It is hard to evaluate the Merced initiative. Judging by

the information that has emerged, the Merced team has not been able to come up with any major architectural innovation. If what I have just said is correct, this is not surprising. It may be that there is unexpected mileage to be obtained from old ideas such as VLIW (very long instruction word) and predicated execution. A critical factor will be the ability of innovative compiler techniques to expose instruction-level parallelism on a much greater scale than has hitherto been possible. This would be a notable breakthrough. There is an obvious danger that the combined hardware/software complexity of the system will defeat the ends aimed at.

I suspect that we may be seeing a natural swing of the pendulum. Old concerns, such as instruction-set compatibility, are reasserting themselves and the breed fanciers are making their voices heard again. I think it probable that we shall see less attention paid to small factors in processor speed. This will provide a favorable climate for processors with instruction sets compatible with that of Pentium II to make further inroads into the workstation market. However, I do not see RISC workstations being driven out altogether.

Performance of Software Systems

I would like to feel that in the future more attention will be paid to system performance as distinct from processor performance. For this purpose, we need effective benchmarks for software systems.

The benchmarking community is already devoting part of its effort to software benchmarks. I observed this recently when I had the privilege of sitting in on a day's discussions in the SPEC Open Systems Group. In hardware, the breakthrough came with the realization that it was strongly in the interests of competing vendors to have agreed upon benchmarks that were as fair and as proof against cheating as human ingenuity could make them. Once the rules of a contest are defined and seen to be fair, tension is reduced. This is well illustrated by the age-old procedure for dividing an apple between two people: one cuts and the other chooses. In the processor field, a common interest made it possible for otherwise competing companies to collaborate harmoniously in the development of benchmarks that could form the basis for subsequent competition.

The same common interest is there in the case of software benchmarks, but the technical problems are much greater. It is easy to be over-optimistic. Performance measurement thrives where there is competition, and languishes where there is monopoly. Perhaps the greatest hope lies in new areas, such as Java, where competition is fierce. Already, SPEC has a Java-client performance benchmark under active development.

The hardware side of the computer industry has had a severe shake-up in recent years. I would not like to see the software side suffer a similar fate, but a milder shake-up would undoubtedly be for the general good. M

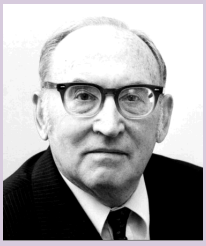

Maurice Wilkes, born in 1913, was for many years head of the Computer Laboratory of the University of Cambridge, where he designed and constructed the EDSAC computer. He was a pioneer of programming for stored-program computers and of microprogramming. He wrote the first paper on cache memories and was an early worker in wide-bandwidth local-area networks.

Now with Olivetti Research, Wilkes is a Distinguished Fellow of the British Computer Society and a Fellow of the Royal Society and of the Royal Academy of Engineering. He is a Foreign Assoicate of the U.S. National Academy of Science and of the National Academy of Engineering.

In 1967, Wilkes delivered the ACM Turing Lecture. He was the recipient of the 1980 Eckert-Mauchly Award and the 1992 Kyoto Prize for Advanced Technology. He has written a number of books, including the first book on computer programming, published in 1951.

For more biographical information, access the Web at *www.cs.stevens-tech.edu/~nar/614/Wilkes.html*.

For more information on modern high-performance CMOS processors, see the author's book *Computing Perspectives* (Morgan-Kauffmann, 1995).

Most Significant Bits

Continued from page 5

by redesigning the cache and MMU for Cypress's original 601 SPARC CPU (see MPR 1/89, p. 10). The company then developed a new processor, originally known as Pinnacle (see MPR 3/25/92, p. 15) and later as HyperSparc.

At the start, Ross hoped to exploit the growing market for SPARC processors, which were used by several companies. But the biggest SPARC vendor, Sun, continually spurned Ross's designs for its own in-house processors. For a brief period, Sun sold a few HyperSparc machines (see MPR 11/14/94, p. 4), as its own SuperSparc had fallen behind in performance. Once Sun moved to UltraSparc, however, it dropped the Ross chip, leaving HyperSparc to address a dwindling market.

Ross began as a subsidiary of Cypress, but Fujitsu purchased a majority stake in 1993. When the Japanese company was flush with cash, it could afford to cover both Ross and Hal Computer, another SPARC chip house. Now facing its own financial problems, Fujitsu is no longer willing to fund Ross's losses. Ross is hoping to find a buyer, but prospects are bleak, even with the company's stock trading for just $^{1}/_{16}$, or $6^{1}/_{4}$ cents a share. Unless it finds a buyer, the company plans to shut down by the end of the year. —*L.G.*

Continued on page 22

Samsung Forms Alpha Processor Inc.

As PowerPC fades, Alpha is stepping to the plate to challenge Intel in the PC market. With Compaq's acquisition of Digital (see MPR 2/16/98, p. 4) now officially complete, Samsung has taken up the Alpha torch. To bolster its presence in the market, the Korean company has formed a U.S. subsidiary, Alpha Processor Inc., to sell Alpha chips to PC makers.

Compaq/Digital will continue to own and maintain the Alpha architecture. That company will also continue to design new Alpha cores, such as the future 21364. Although Samsung has the right to design its own Alpha chips, Alpha Processor will not undertake any chip design work, at least initially.

This leaves the point of the new company unclear. Samsung believes it has created a new business model, but Alpha Processor appears to be simply a U.S. marketing arm for Samsung. The key challenge for Samsung is to find more buyers for its Alpha products; Alpha Processor is merely a new brand for the same chips. —*L.G.*

Equator Sees Light

After years of secrecy, Equator Technologies *(www.equator.com)*, the media-processor startup founded by Multiflow cofounder John O'Donnell (see MPR 6/2/97, p. 4), has finally identified its corporate partners and its plans for the VLIW processors it is developing.

Equator is codeveloping its architecture with Hitachi's Information Systems Development Lab in Japan. Hitachi will also fab parts based on the architecture and contribute "all-format decoding" digital-TV software. While Equator has not named the new architecture or specified details of the first parts, it plans to target the digital-TV and 3D arcade-game markets.

Equator hopes to distinguish itself from other media-processor vendors by providing highly efficient C-language compilers for its processors, allowing developers to realize the full potential of the new architecture without manually optimizing assembly language. TriMedia provides a basic C compiler for its TM-1 media processor (see MPR 11/13/95, p. 22), but O'Donnell's background in compiler design for Multiflow's VLIW machines suggests Equator's compiler may be significantly more powerful, giving it a unique advantage over its competition. —*P.N.G.*

VM Labs Reveals Project X Media Processor

VM Labs *(www.vmlabs.com)* has revealed new details about its own media-processor plans. Like Equator's design, the company's "Project X" architecture—no relation to Micron's Socket X (see MPR 6/1/98, p. 4)—is intended for applications in consumer electronics.

VM Labs plans to replace the hardwired MPEG-2 decoding logic in DVD players and digital satellite receivers with its Project X media processor. The Project X chip will perform these functions in software, and it can also run 3D games and other entertainment applications. By merging the DVD player and the game platform, VM Labs hopes to achieve an even higher penetration of the consumer-electronics market than either device could individually.

VM Labs has allied itself with Motorola's Semiconductor Product Group, which has a minority investment in the company and will fab the chips. Thomson Consumer Electronics (the maker of RCA, GE, and ProScan products) and Toshiba's consumer-electronics group say they will use the Project X chip.

Eleven developers of entertainment software—including major names such as Activision, Capcom, and Psygnosis—have signed up to develop titles for Project X. This show of hardware and software support seems to give VM Labs a leg up on Equator, but shipments of both product families are at least a year away, and much can happen between now and then. —*P.N.G.*

Conspiracy Speeds 3D Geometry

Plugging a small gap in its x86 product line, IBM Microelectronics has teamed with Rendition, Fujitsu, and Hercules Technology to produce Conspiracy, the first mainstream 3D card with hardware geometry acceleration. Since IBM will not offer a processor with 3DNow (see MPR 6/1/98, p. 18) this year, it needed another way to address the CPU bottleneck on 3D transform and lighting calculations.

At about $15, Fujitsu's FGX-1 is the only geometry processor priced for mainstream PC buyers. IBM worked with Fujitsu and Rendition to define a single-board solution combining geometry acceleration and rendering acceleration in the form of Rendition's V2200 2D/3D engine (see MPR 6/23/97, p. 1). Hercules will manufacture the unique card, which is expected to retail this month for about $149.

Unfortunately, there are serious problems with this plan. Today, Quake and other OpenGL-based titles are the only 3D games that can take immediate advantage of hardware geometry acceleration, which is not supported in Microsoft's Direct3D (D3D). While Rendition and Fujitsu have developed ways to patch D3D applications to use the FGX-1, this technique is not supported by Microsoft and will not help most D3D titles. Geometry acceleration is also not supported by 3Dfx's Glide API, which is used by many 3D games to drive 3Dfx's popular Voodoo adapters.

More important, the FGX-1 performs geometry operations only about as fast as a 266-MHz Pentium II processor, so the new Hercules card will be of interest only to end users with older Socket 7 systems—most particularly, those that cannot be upgraded to use the latest Socket 7 processors. For more modern Socket 7 motherboards, an AMD K6-2 or IDT WinChip 2 3D processor would probably match the 3D performance of the FGX-1 and would provide a substantial speedup for non-3D applications as well.

Geometry acceleration is likely to become an important part of the PC 3D market in 1999, when Microsoft provides the necessary hooks in Direct3D, but the software compatibility and performance of Fujitsu's FGX-1 is clearly inadequate for most of today's users. —*P.N.G.*

SLATER PERSPECTIVE

Coming to Grips with Antitrust

Settlements Likely on Business Practices, But Not on Product Integration

Since I last wrote about the issue of Intel's and Microsoft's power a mere six weeks ago (see MPR 5/11/98, p. 15), a lot has happened: Microsoft negotiated until the 11th hour but ultimately decided to fight the U.S.Department of Justice (DOJ) in court, and the U.S. Federal Trade Commission (FTC) has filed suit against Intel (see MPR 6/22/98, p. 8).

Intel and Microsoft are both aggressive companies, to say the least—and their shareholders have reaped the rewards. Both are also smart companies with legions of lawyers working for them, and it seems very unlikely that either company would knowingly pursue an illegal strategy. But the line between what is illegal and what is merely aggressive is not precisely defined, and it is now apparent that the government's view of where the line belongs is quite different than Intel's and Microsoft's views.

Antitrust concerns have been raised about Intel's business practices for years. There have been prior investigations by the FTC, and there have been lawsuits by AMD and Cyrix. Until now, nothing has gone very far. One problem faced by attorneys working on the prior cases is that it was difficult—perhaps impossible—to get PC makers to testify about their dealings with Intel. Digital and Intergraph, however, both reached the point where their executives felt the need to act, and the companies' lawsuits against Intel put in the public domain the first on-the-record complaints about Intel.

The Intergraph case (see MPR 5/11/98, p. 16) is especially relevant, since it deals with Intel's practice of withdrawing its nondisclosure agreements from companies with which it is involved in legal disputes. In the preliminary injunction ruling in this case, Judge Edwin Nelson issued the first legal opinion that Intel "has monopoly power in the relevant market of high-performance CPUs," setting the stage for further actions (though Intel has appealed this decision).

Since it is not possible to be a competitive PC vendor without being able to plan in advance for Intel's future processors, as my colleague Linley Gwennap pointed out last December (see MPR 12/29/97, p. 3), Intel can, by deciding who gets advance information, decide who gets to play and who doesn't—an extraordinary power.

The way Intel uses its NDAs clearly would be legal for an ordinary company but is another issue entirely for a company in Intel's position. Intel asserts that even if it has a monopoly—which it disputes—there is no legal basis for requiring it to disclose confidential information to any company, but the government obviously disagrees.

Many PC makers are concerned about the effects their actions might have on their relationship with Intel, which makes it harder for Intel's competitors to gain design wins. Intel says it would not cancel NDAs in retaliation for a customer's use of a non-Intel chip, and its track record supports this claim—but the mere possibility must cause some concern.

In Microsoft's case, the Department of Justice is going far beyond business practices: it is trying to put limits on product definition. As a result, it is understandable that Microsoft felt the need to fight; fundamental product-design issues are at stake. The browser issue has overshadowed issues about Microsoft's business practices, which are the more appropriate subject for DOJ intervention.

Intel and the FTC may be able to reach a negotiated settlement of the current complaint; it seems that Intel could comply with the FTC's demands without any material affect on its business. Intel is not willing to accept the "monopolist" characterization, however, and Intel's executives are no doubt concerned that any concession would be the beginning of a long series of actions—give 'em an inch and they'll take a mile. Intel may well be concerned that other areas, such as how it allocates chips and comarketing funds, will come under scrutiny. Despite this concern, however, the distraction and resource drain of a long, drawn-out fight with the government—which still could end with Intel losing—could be even more damaging.

Should the FTC challenge Intel's right to integrate graphics into its system logic (as it plans to do with the forthcoming Whitney chip set), or whether Intel can integrate these functions on the processor, then Intel would have to fight to the final appeal—just as Microsoft will fight over browser integration. It is appropriate that the government take a close look at the business practices of these two companies, which have an extraordinary degree of control over a critical industry. But having the government involved in product definition—and asking companies to buck natural technology trends—in untenable.

Intel realized what it meant to be a consumer company only when it was faced with a consumer revolt—and a massive write-off—over the FDIV bug (see MPR 1/23/95, p. 4), and the company's attitudes changed profoundly. Similarly, Intel may have another epiphany as a result of the FTC investigation changing its business practices to those appropriate to a company with its immense market power. M

See www.MDRonline.com/slater/antitrust *for more on this subject. I welcome your feedback at* mslater@zd.com.

LITERATURE WATCH

AUDIO/VIDEO

Maximizing AGP performance. AGP is one of the key enabling technologies for high-performance 3D graphics applications running on desktop PCs. Jim Chu and Frank Hady, Intel; *RTC*, 3/98, p. 83, 6 pp.

Shrinking hardware for MPEG-2. The growing trend: integrating video decoding with other embedded functions. Mike Elphick, *Computer Design*, 5/98, p. 88, 3 pp.

BUSES

Low voltage differential signaling reports for bus duty. Based on point-to-point LVDS technology, bus devices combine speed, low power, and low noise. Jeff Child, *Electronic Design*, 5/98, p. 39, 3 pp.

DEVELOPMENT TOOLS

DSP tools get visual to help developers understand their code. As DSP-based applications move into the mainstream, tools are emerging that use graphics and visualization to help harness more of a DSP's power. Tom Williams, *Embedded Systems Development*, 4/3/98, p. 18, 5 pp.

DSP

Focus report: digital signal processors. Fierce competition across an expanding field of applications has DSP vendors fine-tuning features, consolidating functions, and otherwise pushing the price/performance envelope for 16-bit fixed-point units. Gil Bassak, *Integrated System Design*, 5/98, p. 52, 9 pp.

Digital signal processors look to displace microcontrollers. A DSP-only solution can reduce cost and development time over a DSP-microcontroller combo, if the application fits. Richard Nass, *Embedded Systems Development*, 4/3/98, p. 23, 3 pp.

Higher throughput DSP chips take on complex applications. Consumer, telecommunications, and instrumentation gain from DSP chips with reduced costs and lower power consumption. Dave Bursky, *Electronic Design*, 5/98, p. 80, 4 pp.

IC DESIGN

IC design on the World Wide Web. As university researchers tackle the challenges of circuit design over the Web on geographically distributed, heterogeneous tool sets, electronic-design companies are refining Web-based customer services, and Web-based design-management tools are surfacing. Linda Geppert, *IEEE Spectrum*, 6/98, p. 45, 6 pp.

Circuit density and smaller geometries drive CICC's digital sessions. System chips, programmable logic advances, and deep-submicron processes continue to challenge design engineers. Dave Bursky, *Electronic Design*, 5/25/98, p. 40, 4 pp.

Improved topologies, tools make mixed-signal ASICs possible. Targeted system-on-a-chip devices promise better performance, greater security. Charles H. Small, *Computer Design*, 5/98, p. 27, 4 pp.

Optimized ADCs pack resolution, speed and bandwidth on-chip. Suppliers are tweaking finer processes for better matching of sampling capacitors and switching transistors at low voltages. Ashok Bindra, *Electronic Design*, 5/13/98, p. 46, 5 pp.

MEMORY

Optimize memory subsystem for top performance. A better understanding of memory accesses allows DSP memory subsystems to be better matched to the DSP chips. Richard Jaenicke and Paul Taddonio, Sky Computers; *Electronic Design*, 5/25/98, p. 90, 4 pp.

Putting the squeeze on flash memory. Advances in architecture, processing, and packaging allow more flash storage to be crammed into less space. Mike Elphick, *Computer Design*, 5/98, p. 65, 7 pp.

MISCELLANEOUS

FED up with fat tubes. Field emission displays offer the best and brightest of two display worlds: the bright picture of bulky cathode-ray tube devices and the trim flat-panel picture of liquid-crystal displays. Babu R. Chalamala, Motorola, et al; *IEEE Spectrum*, 4/98, p. 42, 10 pp.

Pocket computers ignite OS. Not content with dominating the market for desktop computer operating systems, Microsoft is now involved in a struggle for what could be an even bigger business—the real-time system software required by handhelds, set-top devices, and car PCs. Richard Comerford, *IEEE Spectrum*, 5/98, p. 43, 6 pp.

PERIPHERAL CIRCUITS

Integrated low-voltage ADCs achieve high speed and accuracy. A family of high-speed ADCs taps Δ-Σ modulation and digital filtering to obtain true 16-bit performance at 2.7 V. Ashok Bindra, *Electronic Design*, 5/1/98, p. 67, 3 pp.

PROCESSORS

Adapting to bigger, faster embedded RISC. No longer just an industrial phenomenon, embedded RISC is riding high on the new wave of consumer products. Jim Turley, MicroDesign Resources; *Computer Design*, 5/98, p. 81, 4 pp.

Information appliances: from Web phones to smart refrigerators. Embedded information processors are spawning many unexpected applications as they use LAN and Internet protocols to communicate across almost any network. Lee Goldberg, *Electronic Design*, 3/23/98, p. 69, 8 pp.

High-integration controller tackles automotive and industrial needs. Based on PowerPC, this RISC CPU packs 448 Kbytes of flash and is a complete data-acquisition and control system. Dave Bursky, *Electronic Design*, 4/6/98, p. 40, 5 pp.

SYSTEM DESIGN

Design embedded systems for low power. Advances in IC technology make it possible to create practical low-power designs with relative ease. Brian Kurkoski, *Circuit Cellar*, 6/98, p. 26, 7 pp.

RECENT IC ANNOUNCEMENTS

PART NUMBER	VENDOR	DESCRIPTION	PRICE/QUANTITY	AVAILABILITY
MICROPROCESSOR				
IS89C52	ISSI 408.588.0800	Microcontroller has 8051 core with 8K of on-chip flash memory, 40-MHz clock rate, 32 programmable I/O pins, UART, and three timers.	$4.50/1,000	Prod.—Now
PIC17C752	Microchip 602.786.7668	Microcontroller has 8-bit core, 16K of one-time-programmable code storage, 678 bytes of data RAM, and 12-channel, 10-bit ADC.	$7.89/10,000	Samples—Now Prod.—3Q98
H8S/2144	Hitachi 800.285.1601	Microcontroller has 16-bit CPU, 128K of flash memory, 20-MHz clock frequency, and single 5-V supply; in 80-lead PQFP package.	$15.50/25,000	Prod.—Now
INTERFACE				
CS8952	Cirrus Logic 512.912.3086	Ethernet transceiver handles 10/100-Mbps transmissions, operates from a single 5-V supply, and consumes 135 mA in normal operation.	$8/1,000	Prod.—Now
CSP1152A	Lucent 800.372.2447	Analog-to-digital converter includes dither function for use in wireless base stations; reduces extraneous signals up to 100 dB below max levels.	$35/50,000	Samples—Now Prod.—4Q98
ST75C530	SGS-Thomson 781.259.0300	Monolithic device for telephones, answering machines, and faxes includes DSP and analog front-end functions for V.17 and V.32bis fax/modem.	$12/100,000	Prod.—Now
CS5529	Crystal 512.912.3736	Delta-signal analog-to-digital converter uses a single power supply, handles industrial temperature range, and settles in one conversion cycle.	$3.80/1,000	Prod.—Now
INA121	Burr-Brown 800.548.6132	FET-input instrumentation amplifier has low, 4-pA input-bias current and supply range from 2.3 V to 18 V; in DIP and SO packages.	$2.35/1,000	Prod.—Now
ADS7843	Burr-Brown 800.548.6132	Single-chip touchscreen controller has 12-bit ADC, consumes 0.6 mW at 75 kHz; SSI/SPI interface and "touch interrupt" alert signal.	$4.95/1,000	Prod.—Now
MEMORY				
HM5225xx5A	Hitachi 800.285.1601	Synchronous DRAMs have 256-Mbit capacity with standard packaging for easy design of PC-100 DIMMs; in ×4, ×8, and ×16 organization.	$500/1	Samples—3Q98 Prod.—4Q98
70V9079	IDT 800.345.7015	Synchronous dual-port SRAM has 3.3-V supply voltage, 50-MHz bus interface, and 32K×8 configuration; 5-V version runs at 66 MHz.	$23.80/10,000	Prod.—Now
HN5x2402	Hitachi 800.285.1601	Serial E^2PROMs have 400-kHz two-wire interfaces, 2-Kbit to 64-Kbit capacity, and wide supply range from 1.8 V to 5.5 V.	$0.40/10,000	Samples—Now Prod.—3Q98
M40Z111	SGS-Thomson 781.259.0300	Nonvolatile SRAM has snap-on battery over surface-mounted SRAM, which allows reflow soldering of SRAM and later attachment of battery.	$7.47/1,000	Prod.—Now
MISCELLANEOUS				
SR05	Semtech 805.498.2111	Newest member of RailClamp family of protection diodes arrays comes in small SOT-143 package, with four surge-rated steering diodes for ESD.	$0.94/1,000	Prod.—Now
CS4912	Crystal 512.912.3559	Digital audio device includes 24-bit DSP, RAM, PLL, digital-audio transmitter port, and CD-quality DAC; in 44-lead PLCC package.	$7/100,000	Prod.—Now
IMP5241	IMP 408.434.1467	Terminator IC for low-voltage differential (LVD) SCSI buses provides autoselectable differential or single-ended termination.	$3.56/1,000	Prod.—Now
PROGRAMMABLE LOGIC				
Ultra37032	Cypress 800.858.1810	CPLD is one of a family of Ultra37000 devices with 5-ns propagation delay and in-system reprogrammability; with 32 macrocells.	$1.25/10,000	Prod.—Now
EPF10K30A	Altera 408.544.7000	High-end PLD has 1,728 logic elements, 12,288 bits of SRAM; suitable for 66-MHz PCI interface or 8-bit, 16-tap FIR filter.	$15/25,000	Samples—Now Prod.—4Q98

PATENT WATCH

by Rich Belgard, Contributing Editor

The following U.S. patents related to microprocessors were issued recently. Please send e-mail to belgard@umunhum.stanford.edu *with comments or questions.*

5,701,508

Executing different instructions that cause different data-type operations to be performed on single logical register file
Issued: December 23, 1997
Inventors: Andrew F. Glew, et al
Assignee: Intel
Filed: December 19, 1995
Claims: 35
Methods for executing different instruction sets (e.g., MMX and floating point) on a processor in a register file. A data processor executes a first set of instructions of a first instruction type (e.g., MMX) on what at least logically appears as a single logical register file. To these instructions, the single logical register file appears to be a flat register file. Additionally, the processor executes a second instruction type (e.g., floating point) using the logical register file, where the register file is treated as a stack.

5,701,442

Method of modifying an instruction-set architecture of a computer processor to maintain backward compatibility
Issued: December 23, 1997
Inventors: Ronny Ronen
Assignee: Intel
Filed: September 19, 1995
Claims: 13
A processor architecture's instruction set has preallocated "hintable" NOP instructions. These NOP instructions have no architectural significance. In current-generation processors, these NOPs may have no function. In future-generation processors, however, these "hintable" NOPs may be used, in a backward-compatible manner, to provide microarchitectural hints, such as caching or power hints, to the processor.

5,701,425

Data processor with functional register and data-processing method
Issued: December 23, 1997
Inventor: Shigeru Nakahara
Assignee: Hitachi
Filed: August 28, 1996
Claims: 24
A processor, with instructions having an opcode and an operand, has general-purpose and functional registers. The functional registers have preassigned functions associated with them (e.g., logical NOT), such that the predefined function of a register is performed when data is written to or read from them.

5,699,536

Computer processing system employing dynamic instruction formatting
Issued: December 16, 1997
Inventors: Martin Edward Hopkins, et al
Assignee: IBM
Filed: April 13, 1995
Claims: 46
A processor with a method of superscalar issue, internally similar to VLIW, whereby instructions of the instruction set may be sequentially scheduled to an execution unit. In parallel with the sequential scheduler, instructions are fetched and group-formatted in the processor into "long decoded instructions" (LDIs). Each LDI is issued to multiple second execution units in parallel.

5,696,958

Method and apparatus for reducing delays following the execution of a branch instruction in an instruction pipeline
Issued: December 9, 1997
Inventors: Todd C. Mowry, et al
Assignee: Silicon Graphics
Filed: March 15, 1995
Claims: 20
A pipelined processor with conditional branch instructions. For every branch instruction, the instruction-fetch stage begins fetching both the branch target and the next sequential instruction before the execution unit completes processing the branch instruction.

5,696,956

Dynamically programmable reduced-instruction-set computer with programmable processor loading on program-number-field and program-number-register contents
Issued: December 9, 1997
Inventors: Rahul Razdan, et al
Assignee: Digital
Filed:November 9, 1995
Claims: 6
A PRISC computer with two types of instructions is claimed. One type is a standard RISC type. The other type is specialized for specific applications. The second instruction type comprises a sequence of RISC-like instructions that are precompiled into single instructions, which are then loaded into a programmable functional unit for execution in a single cycle. M

CHART WATCH: MOBILE PROCESSORS

	Digital 21164	IBM P2SC	PowerPC 750	PowerPC 604e	Sun Ultra-2	Sun Ultra-2i	HP PA-8200	HP PA-7300LC	MIPS R10000	Intel Pentium II
Clock rate	**600 MHz**	160 MHz	300 MHz	350 MHz	360 MHz	300 MHz	236 MHz	160 MHz	250 MHz	400 MHz
Cache size	8K/8K/96K	**32K/128K**	32K/32K	32K/32K	16K/16K	16K/16K	None	64K/64K	32K/32K	16K/16K
Issue rate	4 issue	**6 issue**	3 issue	4 issue	4 issue	4 issue	4 issue	2 issue	4 issue	3 x86 instr
Pipe stages	7 stages	5 stages	6 stages	6 stages	6/9 stages	6/9 stages	7–9 stages	5 stages	5–7 stages	**12–14**
Out of order	6 loads	5 instr	5 instr	16 instr	None	None	**56 instr**	None	32 instr	40 ROPs
Rename regs	None	22 fp	6 int/6 fp	12 int/8 fp	None	None	56 total	None	**32/32**	40 total
BHT entries	**2K × 2-bit**	None	512 × 2-bit	512 × 2-bit	512 × 2-bit	512 × 2-bit	1K × 2-bit	None	512 × 2-bit	≥512
TLB entries	48 I/64 D	64 I/64 D	**128/128**	**128/128**	64 I/64 D	64 I/64 D	120 unified	96 unified	64 unified	32 I/64 D
Memory b/w	~400 MB/s	**2.2 GB/s**	~180 MB/s	~180 MB/s	1.3 GB/s	600 MB/s	768 MB/s	213 MB/s	539 MB/s	528 MB/s
Package	CPGA-499	**SCC-1,088**	CBGA-360	CBGA-255	PBGA-521	PBGA-587	LGA-1,085	CPGA-464	CPGA-527	PBGA-528
IC process	0.35µ 4M	0.25µ 5M	0.25µ 5M	0.25µ 5M	0.29µ 4M	0.29µ 4M	0.5µ 4M	0.5µ 4M	0.25µ 4M	0.25µ 4M
Die size	209 mm^2	255 mm^2	67 mm^2	47 mm^2	126 mm^2	150 mm^2	**345 mm^2**	259 mm^2	197 mm^2	131 mm^2
Transistors	9.3 million	**15 million**	6.4 million	5.1 million	3.8 million	4.1 million	3.9 million	9.2 million	6.8 million	7.5 million
Est mfg cost*	$125	**$290**	$40	$30	$70	$85	$270	$95	$130	$65
Power (max)	25 W	30 W	5 W	7 W	20 W	38 W	**>40 W**	15 W	16 W	15 W*
SPEC95b†	**19.2/26.6**	7.1/23.6	13.2/7.76	14.0/12.1	12.8/21.9	9.6/12.0	16.4/25.3	7.3/7.4	14.1/22.6	15.8/11.4
Availability	2Q97	3Q96	3Q97	3Q97	2Q97	1Q98	3Q97	3Q96	3Q97	1Q98
1K list price	Not public	Not public	$495	$645	$3,635§	$470	Not public	Not public	Not public	$722‡

†SPEC95 baseline (int/FP) ‡includes 512K L2 cache §includes 2M L2 cache (Source: vendors except *MDR estimates)

The table above gives the vital statistics for the key RISC processors available today. The table below provides the best reported SPEC95 results for each shipping processor. (Note that some vendors have submitted only SPECint95 or SPECfp95.) The graph compares transistor counts for these devices for the logic (noncache) portion and the complete design.

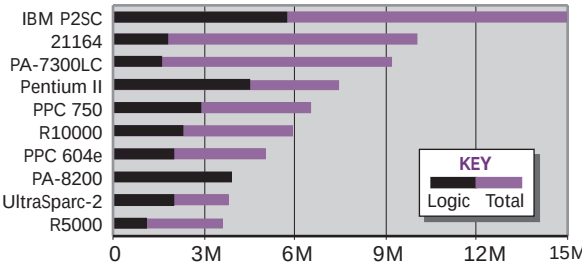

Processor	Digital 21164	HP PA-8200	Intel Pentium II	MIPS R10000	PowerPC 604e	Sun UltraSparc	Sun Ultra-2i	HP 7300LC	Intel P55C	IBM P2SC
System	AlphaServ. 7310 6600	HP9000 K580	SE440BX m'board	SGI Origin2000	RS/6000 H50	Sun Ultra Mod 1360	"Darwin" Ultra 10	HP Visual. Mod B160L	Intel LT430TX	IBM 397 RS/6000
Clock rate	600 MHz	240 MHz	400 MHz	250 MHz	332 MHz	360 MHz	300 MHz	160 MHz	233 MHz	160 MHz
Ext. cache	2M	4M	512K	4M	256K	2M	512K	1M	512K	none
099.go	19.7	17.3	15.1	14.3	17.0	14.4	10.5	9.80	7.80	8.00&
124.m88Ksim	20.9	17.5	15.7	11.8	17.8	12.7	10.1	6.82	8.41	4.67&
126.gcc	20.7	14.8	15.2	13.0	12.8	13.8	9.38	7.16	7.13	6.20&
129.compress	17.5	16.2	12.8	14.9	9.63	14.0	11.3	6.24	4.84	5.55&
130.li	17.3	16.4	16.2	12.2	12.7	10.4	8.38	7.43	8.02	5.21&
132.ijpeg	19.4	14.1	15.7	12.3	16.8	13.0	10.5	5.81	4.66	7.44&
134.perl	20.0	14.0	17.9	16.7	14.6	12.3	8.08	8.08	10.3	5.33&
147.vortex	18.5	22.0	18.2	19.6	12.5	12.5	8.81	7.91	7.55	5.40&
SPECint95b*	**19.2**	**16.4**	**15.8**	**14.1**	**14.0**	**12.8**	**9.57**	**7.32**	**7.12**	**5.88&**
101.tomcatv	34.9	39.5	15.8	28.8	15.1	31.6	17.3	11.1	7.36	46.9
102.swim	49.1	31.2	23.2	42.1	24.4	44.7	23.6	17.9	8.29	56.5
103.su2cor	15.6	17.7	7.53	14.1	5.68	14.9	6.13	3.66	3.05	10.5
104.hydro2d	19.5	14.3	7.25	15.6	6.10	14.9	6.32	3.86	2.95	12.9
107.mgrid	37.3	21.2	7.74	22.9	9.60	24.3	12.7	5.64	2.15	22.8
110.applu	9.4	18.7	7.28	14.4	8.05	12.7	8.21	4.99	2.01	23.2
125.turb3d	25.9	19.1	10.7	17.7	15.0	17.3	12.0	6.70	4.69	22.7
141.apsi	26.2	28.5	14.3	20.1	9.22	26.0	12.2	7.69	4.81	11.5
145.fpppp	43.2	42.6	17.7	35.2	36.9	22.8	16.6	12.4	5.65	35.1
146.wave5	33.0	37.7	11.8	32.0	12.8	25.9	16.0	9.54	6.03	30.1
SPECfp95b*	**26.6**	**25.3**	**11.4**	**22.6**	**12.1**	**21.9**	**12.0**	**7.38**	**4.23**	**23.6**

*SPEC95 baseline ‡SGI Indy R5000 &IBM 595 w/P2SC-135 (Source: SPEC)

Remember Your Anatomy Lessons?

Bruce Shriver and Bennett Smith do. They've authored *The Anatomy of a High-Performance Microprocessor: A Systems Perspective*. Published by the IEEE Computer Society, the 584-page book (which comes with a CD-ROM) examines microarchitectural details in designing a high-performance CPU, including "platform and system issues." AMD's K6-3D acts as cadaver.

The book (ISBN 0-8186-8400-3) and access to its associated Web site (*www.computer.org/books/anatomy*) cost $44 for nonmembers. More information can be found at *www.computer.org/cspress/CATALOG/anatomy.htm*.

Giving Up PCs for PLCs?

Then you might want to check out *Industrial Controls Intelligence & the PLC Insiders' Newsletter*, a monthly subscriber-supported newsletter put out by Jack Grenard and Carefree Communications. Annual subscriptions run $195. For more information, contact Carefree (Carefree, Ariz.) at 602.488.1462 or write *jgrenard@aol.com*.

Gartner Group on Semiconductors

Coincident with Semicon/West, Gartner Group consultants will be discussing semiconductor manufacturing issues July 15 at San Francisco's Palace Hotel. For more information, call 800.645.6395 or visit *www.winntech.net/gartner*.

You Still Have Time to Implement

Such is the comforting claim of *Year 2000: Best Practices for the Millennium*, a 660-page book edited by Dick Lefkon and published by Prentice-Hall. The book includes well over 100 articles and chapters contributed by representatives of NIST, Ernst & Young, the U.S. Department of Defense, IBM, the State of California, and, in one case, the U.S. Congress.The $40 book (ISBN 0-13-646506-4) is available through the publisher at *www.phptr.com*. IEEE members; $49 for nonmembers.

Sequel to a Millennial Best Seller

Ian S. Hayes and William M. Ulrich, self-described authors of the #1 Year 2000 Best Seller, have completed *The Year 2000 Software Crisis: The Continuing Challenge*. This 440-page treatise includes chapters on risk mitigation, legal issues, protections, testing basics, and contingency planning. The $40 book (ISBN 0-13-960154-6) is available through Prentice-Hall at *www.phptr.com*.

Never Be Out of Year 2000 Contact

MBS2000, a London-based consultancy specializing in Y2K issues, has opened what it describes as the first Web site dedicated to millennial bug product compliance. Located at *www.mbs2000.com*, the site lists hardware and software products that are in compliance and provides fixes for those that aren't.

SUBSCRIPTION INFORMATION

To subscribe to *Microprocessor Report*, contact our customer service department by phone, 707.824.4001; fax, 707.823.0504; e-mail, *cs@mdr.zd.com;* or Web, *www.MDRonline.com.*

For European orders, contact Parkway Gordon by phone, 44.1491.875386, or fax, 44.1491.875524; or send e-mail to *parkway@rmplc.co.uk,* or visit the Web at *www.parkway.co.uk.*

	U.S. and Canada*	Europe	Elsewhere
One year	$595	£450	$695
Two years	$1,095	£795	$1,295

**Sales tax applies in the following states: GA, KY, MA, TX, and WA. GST tax applies in Canada.*

Microprocessor Report (ISSN 0899-9341) is published every three weeks, 17 issues per year. Back issues are available on paper and CD-ROM. Volume reprints of individual articles are also available.

Ship to: